



OcNOS®

**Open Compute Network Operating System
for Service Providers**

Multi-Protocol Label Switching

Version 7.0.0

February 2026

©2026 IP Infusion Inc. All Rights Reserved.

This documentation is subject to change without notice. The software described in this document and this documentation are furnished under a license agreement or nondisclosure agreement. The software and documentation may be used or copied only in accordance with the terms of the applicable agreement. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or any means electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's internal use without the written permission of IP Infusion Inc.

IP Infusion Inc.

3979 Freedom Circle, Suite 900

Santa Clara, CA 95054

+1 408-400-1900

<http://www.ipinfusion.com/>

For support, questions, or comments via E-mail, contact:

support@ipinfusion.com

Trademarks:

IP Infusion and OcNOS are trademarks or registered trademarks of IP Infusion. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Use of certain software included in this equipment is subject to the IP Infusion, Inc. End User License Agreement at <http://www.ipinfusion.com/license>. By using the equipment, you accept the terms of the End User License Agreement.

| CONTENTS

| | |
|--|------------|
| Contents | 3 |
| Preface | 103 |
| About this Guide | 103 |
| Audience | 103 |
| Conventions | 103 |
| IP Infusion Product Release Version | 103 |
| Related Documentation | 104 |
| Feature Availability | 104 |
| Migration Guide | 104 |
| IP Maestro Support | 104 |
| Technical Support | 104 |
| Technical Sales | 104 |
| Technical Documentation | 104 |
| Documentation Disclaimer | 105 |
| Comments | 105 |
| Command Line Interface | 106 |
| Overview | 106 |
| Chapter Organization | 106 |
| Command Line Interface Help | 106 |
| Command Completion | 107 |
| Command Abbreviations | 107 |
| Command Line Errors | 108 |
| Command Negation | 108 |
| Syntax Conventions | 108 |
| Variable Placeholders | 109 |
| Command Description Format | 110 |
| Keyboard Operations | 110 |
| Show Command Modifiers | 111 |
| Begin Modifier | 111 |
| Include Modifier | 112 |
| Exclude Modifier | 112 |
| Redirect Modifier | 113 |
| Last Modifier | 113 |
| String Parameters | 113 |
| Command Modes | 114 |
| Command Mode Tree | 115 |
| Transaction-based Command-line Interface | 115 |
| Label Distribution Protocol Configuration | 117 |
| LDP Configuration | 119 |
| Overview | 119 |
| LDP Adjacencies | 119 |

| | |
|---|-----|
| LDP Session | 119 |
| Forwarding Equivalence Class | 119 |
| Label Generation | 119 |
| Label Distribution Modes | 120 |
| Label Retention Mode | 120 |
| LSP Control | 120 |
| Loop Detection | 120 |
| Configure LDP | 120 |
| LDP MD5 Authentication | 126 |
| Validation | 127 |
| PE1 | 127 |
| PE2 | 128 |
| Configure LDP MD5 for Targeted LDP Session | 128 |
| Removing MD5 Authentication for LDP Session | 130 |
| Validation | 131 |
| LDP Session Count | 131 |
| FTN, SWAP, and POP Entries | 131 |
| MPLS LDP PING and TRACEROUTE | 132 |
| LDP Session Protection | 133 |
| Validation | 139 |
| Without session protection enabled | 139 |
| Configure Session Protection | 141 |
| Validation | 142 |
| With session protection command enabled | 142 |
| Link Down Validation | 146 |
| After link down | 146 |
| Up link Validation | 150 |
| LDP Inter-Area Configuration | 155 |
| Overview | 155 |
| Configure LDP Inter-Area | 155 |
| Validation | 161 |
| MPLS LDP-IGP Synchronization | 171 |
| Overview | 171 |
| Prerequisites | 171 |
| Topology | 171 |
| LDP-IGP Synchronization with OSPF | 172 |
| Validation | 174 |
| LDP-IGP Synchronization | 175 |
| PE1 Validation | 176 |
| PE2 Validation | 177 |
| LDP-IGP Synchronization with IS-IS | 177 |
| Validation | 180 |
| LDP-IGP SYNC Configuration | 182 |
| PE1 Validation | 184 |
| PE2 Validation | 184 |
| LDP-FRR Configuration | 186 |
| Topology | 187 |

| | |
|---|-----|
| LDP-FRR with OSPF as IGP Configuration | 187 |
| PE1 | 187 |
| P1 | 190 |
| P2 | 192 |
| P3 | 195 |
| PE2 | 197 |
| Validation | 199 |
| LDP FRR with OSPF Configuration | 234 |
| Validation | 235 |
| LDP-FRR with ISIS as IGP Configuration | 278 |
| PE1 | 278 |
| P1 | 280 |
| P2 | 282 |
| P3 | 285 |
| PE2 | 286 |
| Validation | 288 |
| LDP FRR with IS-IS Configuration | 319 |
| Validation | 320 |
| Backup Path based on Route-Map Prefixes | 334 |
| Validation | 335 |
| TCP MSS configuration for LDP sessions | 351 |
| Overview | 351 |
| Feature Characteristics | 351 |
| Benefits | 352 |
| Prerequisites | 352 |
| Configuration | 352 |
| Enable Label Switching | 352 |
| Topology | 352 |
| Configuration | 352 |
| Validation | 356 |
| Configure TCP MSS on ALL neighbor | 360 |
| Validation | 364 |
| Configuration of TCP MSS with Auto-targeted | 368 |
| Validation | 372 |
| TCP MSS for LDP Session Command | 377 |
| neighbor tcp-mss | 377 |
| Abbreviations | 378 |
| Glossary | 378 |
| LDP Remote Loop-Free Alternate (RLFA) | 379 |
| RLFA Terminology | 379 |
| Repair tunnel | 380 |
| P-space | 380 |
| Extended P-space | 380 |
| Q-space | 380 |
| PQ node | 380 |
| Remote LFA (RLFA) | 380 |
| Establishing RLFA Tunnel | 380 |

| | |
|---|-----|
| IS-IS | 380 |
| LDP | 381 |
| NSM | 381 |
| HSL | 381 |
| Configure LDP Remote Loop-Free Alternate (RLFA) | 381 |
| Assumptions and Limitations | 381 |
| Topology | 382 |
| PE1 Configurations | 382 |
| PE2 Configurations | 383 |
| PE3 Configurations | 385 |
| PE4 Configurations | 386 |
| P1 Configurations | 387 |
| P2 Configurations | 388 |
| Validation | 390 |
| LDP MD5 Password For Auto Targeted Sessions | 416 |
| Configuration | 416 |
| Topology | 416 |
| PE1 | 417 |
| P1 | 418 |
| PE2 | 419 |
| PE3 | 421 |
| P2 | 422 |
| PE4 | 424 |
| Validation | 425 |
| PE1 | 425 |
| LDP ECMP Configuration | 435 |
| Topology | 435 |
| Configurations | 435 |
| Validation | 450 |
| LDP Graceful Restart | 465 |
| Topology | 465 |
| Configuration | 465 |
| Validation | 467 |
| PE1 Before LDP Graceful Restart | 467 |
| PE1 after Graceful Restart | 469 |
| PE1 Restarting the LDP process | 471 |
| P Before Restart LDP | 472 |
| P After LDP Graceful Restart | 474 |
| PE2 Before LDP Graceful Restart | 475 |
| PE2 After LDP Graceful Restart | 477 |
| PE2 Restarting the LDP process | 477 |
| LDP Tunneling over RSVP-TE | 479 |
| Overview | 479 |
| Feature Characteristics | 479 |
| Benefits | 479 |
| Prerequisites | 479 |
| Limitations | 479 |

| | |
|---|------------|
| Configuration for LDP Tunneling Over RSVP | 480 |
| Topology | 480 |
| Configure LDP Tunneling over RSVP | 480 |
| Running Configuration | 484 |
| Validation | 486 |
| Commands for LDP Tunneling over RSVP-TE | 490 |
| ldp-tunneling | 491 |
| prefer-tunnel-in-tunnel rsvp | 492 |
| show ldp tunneling | 493 |
| show ldp tunneling fec | 494 |
| show ldp tunneling-tunnels | 495 |
| Glossary | 495 |
| Revert Timer for Revertive Switchover for LDP Pseudowires | 497 |
| Overview | 497 |
| Characteristics | 497 |
| Benefits | 497 |
| Prerequisites | 497 |
| Configuration | 497 |
| Topology | 497 |
| MPLS LDP Configuration | 498 |
| VPWS Configuration | 501 |
| VPLS Configuration | 502 |
| Reverter Timer Running Configuration | 504 |
| Validation | 508 |
| Revert Timer Commands | 512 |
| Implementation Examples | 512 |
| Troubleshooting Revert Timer for Revertive Switchover Issue | 512 |
| Glossary | 512 |
| Label Distribution Protocol Command Reference | 513 |
| LDP Commands | 516 |
| auth md5 password | 518 |
| Command Syntax | 518 |
| Parameters | 518 |
| Default | 518 |
| Command Mode | 518 |
| Applicability | 518 |
| Example | 518 |
| advertise-labels | 519 |
| Command Syntax | 519 |
| Parameters | 519 |
| Default | 519 |
| Command Mode | 519 |
| Applicability | 519 |
| Examples | 519 |
| advertise-label-for-default-route | 521 |
| Command Syntax | 521 |

| | |
|---|-----|
| Parameters | 521 |
| Default | 521 |
| Command Mode | 521 |
| Applicability | 521 |
| Examples | 521 |
| advertisement-mode | 522 |
| Command Syntax | 522 |
| Default | 522 |
| Command Mode | 522 |
| Applicability | 522 |
| Example | 522 |
| auto-targeted-session | 523 |
| Command Syntax | 523 |
| Parameters | 523 |
| Command Mode | 523 |
| Applicability | 523 |
| Example | 523 |
| clear ldp adjacency | 524 |
| Command Syntax | 524 |
| Parameters | 524 |
| Command Mode | 524 |
| Applicability | 524 |
| Example | 524 |
| clear ldp session | 525 |
| Command Syntax | 525 |
| Parameters | 525 |
| Command Mode | 525 |
| Applicability | 525 |
| Example | 525 |
| clear ldp statistics | 526 |
| Command Syntax | 526 |
| Parameters | 526 |
| Command Mode | 526 |
| Applicability | 526 |
| Example | 526 |
| clear ldp statistics advertise-labels | 527 |
| Command Syntax | 527 |
| Parameters | 527 |
| Command Mode | 527 |
| Applicability | 527 |
| Example | 527 |
| control-mode | 528 |
| Command Syntax | 528 |
| Parameters | 528 |
| Command Mode | 528 |
| Default | 528 |
| Applicability | 529 |

| | |
|----------------------------------|-----|
| Example | 529 |
| debug ldp advertise-labels | 530 |
| Command Syntax | 530 |
| Parameters | 530 |
| Command Mode | 530 |
| Applicability | 530 |
| Example | 530 |
| debug ldp all | 531 |
| Command Syntax | 531 |
| Parameters | 531 |
| Command Mode | 531 |
| Applicability | 531 |
| Example | 531 |
| debug ldp dsm | 532 |
| Command Syntax | 532 |
| Parameters | 532 |
| Command Mode | 532 |
| Applicability | 532 |
| Example | 532 |
| debug ldp events | 533 |
| Command Syntax | 533 |
| Parameters | 533 |
| Command Mode | 533 |
| Applicability | 533 |
| Example | 533 |
| debug ldp fsm | 534 |
| Command Syntax | 534 |
| Parameters | 534 |
| Command Mode | 534 |
| Applicability | 534 |
| Example | 534 |
| debug ldp hexdump | 535 |
| Command Syntax | 535 |
| Parameters | 535 |
| Command Mode | 535 |
| Applicability | 535 |
| Example | 535 |
| debug ldp inter-area | 536 |
| Command Syntax | 536 |
| Parameters | 536 |
| Command Mode | 536 |
| Applicability | 536 |
| Example | 536 |
| debug ldp nsm | 537 |
| Command Syntax | 537 |
| Parameters | 537 |
| Command Mode | 537 |

| | |
|------------------------|-----|
| Applicability | 537 |
| Example | 537 |
| debug ldp packet | 538 |
| Command Syntax | 538 |
| Parameters | 538 |
| Command Mode | 538 |
| Applicability | 538 |
| Example | 538 |
| debug ldp usm | 539 |
| Command Syntax | 539 |
| Parameters | 539 |
| Command Mode | 539 |
| Applicability | 539 |
| Example | 539 |
| debug ldp vc usm | 540 |
| Command Syntax | 540 |
| Parameters | 540 |
| Command Mode | 540 |
| Applicability | 540 |
| Example | 540 |
| disable-ldp | 541 |
| Command Syntax | 541 |
| Parameters | 541 |
| Default | 541 |
| Command Mode | 541 |
| Applicability | 541 |
| Examples | 541 |
| enable-ldp | 542 |
| Command Syntax | 542 |
| Parameters | 542 |
| Default | 542 |
| Command Mode | 542 |
| Applicability | 542 |
| Examples | 542 |
| explicit-null | 543 |
| Command Syntax | 543 |
| Parameters | 543 |
| Default | 543 |
| Command Mode | 543 |
| Applicability | 543 |
| Examples | 543 |
| fast-reroute | 544 |
| Command Syntax | 544 |
| Parameters | 544 |
| Command Mode | 544 |
| Applicability | 544 |
| Examples | 544 |

| | |
|-------------------------------|-----|
| global-merge-capability | 545 |
| Command Syntax | 545 |
| Parameters | 545 |
| Default | 545 |
| Command Mode | 545 |
| Applicability | 545 |
| Examples | 545 |
| graceful-restart | 546 |
| Command Syntax | 546 |
| Parameters | 546 |
| Default | 546 |
| Command Mode | 546 |
| Applicability | 546 |
| Examples | 546 |
| hello-interval | 548 |
| Command Syntax | 548 |
| Parameters | 548 |
| Default | 548 |
| Command Mode | 548 |
| Applicability | 548 |
| Examples | 548 |
| hold-time | 549 |
| Command Syntax | 549 |
| Parameters | 549 |
| Default | 549 |
| Command Mode | 549 |
| Applicability | 549 |
| Example | 549 |
| import-bgp-routes | 550 |
| Command Syntax | 550 |
| Parameters | 550 |
| Default | 550 |
| Command Mode | 550 |
| Applicability | 550 |
| Example | 550 |
| inter-area-lsp | 551 |
| Command Syntax | 551 |
| Parameters | 551 |
| Command Mode | 551 |
| Applicability | 551 |
| Example | 551 |
| keepalive-interval | 552 |
| Command Syntax | 552 |
| Parameters | 552 |
| Default | 552 |
| Command Mode | 552 |
| Applicability | 552 |

| | |
|------------------------------|-----|
| Example | 552 |
| keepalive-timeout | 553 |
| Command Syntax | 553 |
| Parameters | 553 |
| Default | 553 |
| Applicability | 553 |
| Example | 553 |
| label-retention-mode | 554 |
| Command Syntax | 554 |
| Parameters | 555 |
| Default | 555 |
| Command Mode | 555 |
| Applicability | 555 |
| Example | 555 |
| ldp advertisement-mode | 556 |
| Command Syntax | 556 |
| Parameters | 556 |
| Default | 556 |
| Command Mode | 556 |
| Applicability | 556 |
| Example | 556 |
| ldp hello-interval | 557 |
| Command Syntax | 557 |
| Parameters | 557 |
| Default | 557 |
| Command Mode | 557 |
| Applicability | 557 |
| Examples | 557 |
| ldp hold-time | 558 |
| Command Syntax | 558 |
| Parameters | 558 |
| Default | 558 |
| Command Mode | 558 |
| Applicability | 558 |
| Example | 558 |
| ldp keepalive-interval | 559 |
| Command Syntax | 559 |
| Parameters | 559 |
| Default | 559 |
| Command Mode | 559 |
| Applicability | 559 |
| Examples | 559 |
| ldp keepalive-timeout | 560 |
| Command Syntax | 560 |
| Parameters | 560 |
| Default | 560 |
| Command Mode | 560 |

| | |
|-------------------------------------|-----|
| Applicability | 560 |
| Example | 560 |
| ldp label-retention-mode | 561 |
| Command Syntax | 561 |
| Parameters | 561 |
| Default | 561 |
| Command Mode | 561 |
| Applicability | 561 |
| Example | 561 |
| ldp multicast-hellos | 562 |
| Command Syntax | 562 |
| Parameters | 562 |
| Default | 562 |
| Command Mode | 562 |
| Applicability | 562 |
| Example | 562 |
| ldp-optimization | 563 |
| Command Syntax | 563 |
| Parameters | 563 |
| Default | 563 |
| Command Mode | 563 |
| Applicability | 563 |
| Example | 563 |
| loop-detection | 564 |
| Command Syntax | 564 |
| Parameters | 564 |
| Default | 564 |
| Command Mode | 564 |
| Applicability | 564 |
| Example | 564 |
| loop-detection-hop-count | 565 |
| Command Syntax | 565 |
| Parameters | 565 |
| Default | 565 |
| Command Mode | 565 |
| Applicability | 565 |
| Examples | 565 |
| loop-detection-path-vec-count | 566 |
| Command Syntax | 566 |
| Parameters | 566 |
| Default | 566 |
| Command Mode | 566 |
| Applicability | 566 |
| Example | 566 |
| mpls ldp-igp sync-delay | 567 |
| Command Syntax | 567 |
| Parameters | 567 |

| | |
|----------------------------------|-----|
| Default | 567 |
| Command Mode | 567 |
| Applicability | 567 |
| Example | 567 |
| mpls ldp-igp sync isis | 568 |
| Command Syntax | 568 |
| Parameters | 568 |
| Default | 568 |
| Command Mode | 568 |
| Applicability | 568 |
| Example | 568 |
| mpls ldp-igp sync ospf | 569 |
| Command Syntax | 569 |
| Parameters | 569 |
| Default | 569 |
| Command Mode | 569 |
| Applicability | 569 |
| Example | 569 |
| multicast-hellos | 570 |
| Command Syntax | 570 |
| Parameters | 570 |
| Default | 570 |
| Command Mode | 570 |
| Applicability | 570 |
| Example | 570 |
| neighbor | 571 |
| Command Syntax | 571 |
| Parameters | 571 |
| Default | 571 |
| Command Mode | 571 |
| Applicability | 571 |
| Example | 571 |
| neighbor auth md5 password | 572 |
| Command Syntax | 572 |
| Parameters | 572 |
| Default | 572 |
| Command Mode | 572 |
| Applicability | 572 |
| Example | 572 |
| neighbor prefix-list | 573 |
| Command Syntax | 573 |
| Parameters | 573 |
| Default | 573 |
| Command Mode | 573 |
| Applicability | 573 |
| Example | 573 |
| propagate-release | 574 |

| | |
|-----------------------------|-----|
| Command Syntax | 574 |
| Parameters | 574 |
| Default | 574 |
| Command Mode | 574 |
| Applicability | 574 |
| Example | 574 |
| pw-status-tlv | 575 |
| Command Syntax | 575 |
| Parameters | 575 |
| Default | 575 |
| Command Mode | 575 |
| Applicability | 575 |
| Example | 575 |
| request-labels-for | 576 |
| Command Syntax | 576 |
| Parameters | 576 |
| Command Mode | 576 |
| Applicability | 576 |
| Examples | 576 |
| request-retry | 577 |
| Command Syntax | 577 |
| Parameters | 577 |
| Default | 577 |
| Command Mode | 577 |
| Applicability | 577 |
| Examples | 577 |
| request-retry-timeout | 578 |
| Command Syntax | 578 |
| Parameter | 578 |
| Default | 578 |
| Command Mode | 578 |
| Applicability | 578 |
| Example | 578 |
| restart ldp graceful | 579 |
| Command Syntax | 579 |
| Parameter | 579 |
| Command Mode | 579 |
| Applicability | 579 |
| Example | 579 |
| rlfa-ilm-optimization | 580 |
| Command Syntax | 580 |
| Parameter | 580 |
| Default | 580 |
| Command Mode | 580 |
| Applicability | 580 |
| Example | 580 |
| router-id | 581 |

| | |
|------------------------------------|-----|
| Command Syntax | 581 |
| Parameter | 581 |
| Default | 581 |
| Command Mode | 581 |
| Applicability | 581 |
| Example | 581 |
| router ldp | 582 |
| Command Syntax | 582 |
| Parameters | 582 |
| Default | 582 |
| Command Mode | 582 |
| Applicability | 582 |
| Example | 582 |
| session-group name | 583 |
| Command Syntax | 583 |
| Parameters | 583 |
| Default | 583 |
| Command Mode | 583 |
| Applicability | 583 |
| Example | 583 |
| snmp restart ldp | 584 |
| Command Syntax | 584 |
| Parameters | 584 |
| Default | 584 |
| Command Mode | 584 |
| Applicability | 584 |
| Examples | 584 |
| targeted-peer hello-interval | 585 |
| Command Syntax | 585 |
| Parameters | 585 |
| Default | 585 |
| Command Mode | 585 |
| Applicability | 585 |
| Examples | 585 |
| targeted-peer ipv4 | 586 |
| Command Syntax | 586 |
| Parameter | 586 |
| Default | 586 |
| Command Mode | 586 |
| Applicability | 586 |
| Examples | 586 |
| transport-address ipv4 | 587 |
| Command Syntax | 587 |
| Parameters | 587 |
| Default | 587 |
| Command Mode | 587 |
| Applicability | 587 |

| | |
|-------------------------------|-----|
| Example | 587 |
| targeted-peer tunneling | 588 |
| Command Syntax | 588 |
| Parameters | 588 |
| Default | 588 |
| Command Mode | 588 |
| Applicability | 588 |
| Examples | 588 |
| targeted-peer hold-time | 589 |
| Command Syntax | 589 |
| Parameters | 589 |
| Default | 589 |
| Command Mode | 589 |
| Applicability | 589 |
| Examples | 589 |
| targeted-peer-hold-time | 590 |
| Command Syntax | 590 |
| Parameter | 590 |
| Default | 590 |
| Command Mode | 590 |
| Applicability | 590 |
| Example | 590 |
| LDP Show Commands | 591 |
| show ldp label-pool | 592 |
| Command Syntax | 592 |
| Parameter | 592 |
| Command Mode | 592 |
| Applicability | 592 |
| Example | 592 |
| show debugging ldp | 593 |
| Command Syntax | 593 |
| Parameter | 593 |
| Command Mode | 593 |
| Applicability | 593 |
| Example | 593 |
| show ldp | 594 |
| Command Syntax | 594 |
| Parameter | 594 |
| Command Mode | 594 |
| Applicability | 594 |
| Examples | 594 |
| show ldp adjacency | 596 |
| Command Syntax | 596 |
| Parameter | 596 |
| Command Mode | 596 |
| Applicability | 596 |
| Example | 596 |

| | |
|---------------------------------------|-----|
| show ldp adjacency count | 597 |
| Command Syntax | 597 |
| Parameter | 597 |
| Command Mode | 597 |
| Applicability | 597 |
| Example | 597 |
| show ldp advertise-labels | 598 |
| Command Syntax | 598 |
| Parameter | 598 |
| Command Mode | 598 |
| Applicability | 598 |
| Example | 598 |
| show ldp fec | 599 |
| Command Syntax | 599 |
| Parameter | 599 |
| Command Mode | 599 |
| Applicability | 599 |
| Example | 599 |
| show ldp igp sync | 601 |
| Parameter | 601 |
| Command Mode | 601 |
| Applicability | 601 |
| Example | 601 |
| show ldp inter-area-fecs | 602 |
| Command Syntax | 602 |
| Parameter | 602 |
| Command Mode | 602 |
| Applicability | 602 |
| Example | 602 |
| show ldp inter-area-fecs prefix | 603 |
| Command Syntax | 603 |
| Parameter | 603 |
| Command Mode | 603 |
| Applicability | 603 |
| Examples | 603 |
| show ldp interface | 604 |
| Command Syntax | 604 |
| Parameter | 604 |
| Command Mode | 604 |
| Applicability | 604 |
| Examples | 604 |
| show ldp lsp | 606 |
| Command Syntax | 606 |
| Parameter | 606 |
| Command Mode | 606 |
| Applicability | 606 |
| Example | 606 |

| | |
|--|-----|
| show ldp mpls-l2-circuit | 609 |
| Command Syntax | 609 |
| Parameter | 609 |
| Command Mode | 609 |
| Applicability | 609 |
| Example | 609 |
| show ldp rfa-routes | 612 |
| Command Syntax | 612 |
| Parameter | 612 |
| Command Mode | 612 |
| Applicability | 612 |
| Example | 612 |
| show ldp rfa-routes count | 613 |
| Command Syntax | 613 |
| Parameter | 613 |
| Command Mode | 613 |
| Applicability | 613 |
| Example | 613 |
| show ldp routes | 614 |
| Command Syntax | 614 |
| Parameter | 614 |
| Command Mode | 614 |
| Applicability | 614 |
| Example | 614 |
| show ldp routes count | 616 |
| Command Syntax | 616 |
| Parameter | 616 |
| Command Mode | 616 |
| Applicability | 616 |
| Example | 616 |
| show ldp session | 617 |
| Command Syntax | 617 |
| Parameter | 617 |
| Command Mode | 617 |
| Applicability | 617 |
| Example | 617 |
| show ldp session all count | 620 |
| Command Syntax | 620 |
| Parameter | 620 |
| Command Mode | 620 |
| Applicability | 620 |
| Example | 620 |
| show ldp session multicast count | 621 |
| Command Syntax | 621 |
| Parameter | 621 |
| Command Mode | 621 |
| Applicability | 621 |

| | |
|--|-----|
| Example | 621 |
| show ldp session targeted count | 622 |
| Command Syntax | 622 |
| Parameter | 622 |
| Command Mode | 622 |
| Applicability | 622 |
| Example | 622 |
| show ldp statistics | 623 |
| Command Syntax | 623 |
| Parameter | 623 |
| Command Mode | 623 |
| Applicability | 623 |
| Example | 623 |
| show ldp statistics advertise-labels | 625 |
| Command Syntax | 625 |
| Parameter | 625 |
| Command Mode | 625 |
| Applicability | 625 |
| Example | 625 |
| show ldp targeted-peers | 626 |
| Command Syntax | 626 |
| Parameter | 626 |
| Command Mode | 626 |
| Applicability | 626 |
| Example | 626 |
| show ldp upstream | 627 |
| Command Syntax | 627 |
| Parameter | 627 |
| Command Mode | 627 |
| Applicability | 627 |
| Example | 627 |
| show ldp downstream | 629 |
| Command Syntax | 629 |
| Parameter | 629 |
| Command Mode | 629 |
| Applicability | 629 |
| Example | 629 |
| show ldp vpls | 631 |
| Command Syntax | 631 |
| Parameter | 631 |
| Command Mode | 631 |
| Applicability | 631 |
| Example | 631 |
| show mpls ldp discovery | 634 |
| Command Syntax | 634 |
| Parameter | 634 |
| Command Mode | 634 |

| | |
|--|------------|
| Applicability | 634 |
| Example | 634 |
| show mpls ldp neighbor | 635 |
| Command Syntax | 635 |
| Parameter | 635 |
| Command Mode | 635 |
| Applicability | 635 |
| Example | 635 |
| show mpls ldp parameter | 636 |
| Command Syntax | 636 |
| Parameter | 636 |
| Command Mode | 636 |
| Applicability | 636 |
| Example | 636 |
| show ldp tunneling-fec | 638 |
| Command Syntax | 638 |
| Parameters | 638 |
| Default | 638 |
| Command Mode | 638 |
| Applicability | 638 |
| Examples | 638 |
| show ldp tunneling | 639 |
| Command Syntax | 639 |
| Parameters | 639 |
| Default | 639 |
| Command Mode | 639 |
| Applicability | 639 |
| Examples | 639 |
| show ldp tunneling-tunnels | 640 |
| Command Syntax | 640 |
| Parameters | 640 |
| Default | 640 |
| Command Mode | 640 |
| Applicability | 640 |
| Examples | 640 |
| Resource Reservation Protocol Configuration | 641 |
| RSVP-TE Configuration | 643 |
| RSVP-TE Overview | 643 |
| RSVP-TE Architecture | 643 |
| Configure RSVP-TE | 644 |
| Enable Label Switching - Minimal Configuration | 644 |
| Establish a Trunk with CSPF Disabled | 647 |
| Establish a Trunk Using CSPF | 647 |
| Mapping a Route to a Trunk | 648 |
| Establish a Trunk Using Explicitly-Defined Path | 649 |
| Add a Secondary LSP to the Trunk | 650 |

| | |
|--|-----|
| Add Multiple Secondary LSP to the trunk | 652 |
| Add Administrative Group Constraints to an LSP | 672 |
| Configure Global Parameters | 674 |
| MPLS RSVP PING and TRACEROUTE | 675 |
| MPLS RSVP Entropy Label Capabilities | 675 |
| MPLS RSVP LSP Re-optimization | 676 |
| RSVP-TE Facility Backup (Facility Bypass) | 677 |
| Topology | 677 |
| Configuration | 677 |
| PE1 | 677 |
| P1 | 678 |
| P2 | 679 |
| PE2 | 680 |
| RSVP Path on PE1 | 681 |
| Validation | 682 |
| OSPF Neighborhood | 682 |
| RSVP Session | 682 |
| Limitations | 687 |
| Facility Bypass with Ring Topology Configuration | 688 |
| Topology | 689 |
| Configuration | 689 |
| R1 | 689 |
| R2 | 691 |
| R3 | 692 |
| R4 | 693 |
| R5 | 694 |
| Validation | 695 |
| RSVP Session | 695 |
| RSVP Bypass | 702 |
| RSVP-Multipath Configuration | 704 |
| Topology | 704 |
| Configuration | 704 |
| R1 | 704 |
| R2 | 706 |
| R3 | 707 |
| R4 | 708 |
| R5 | 708 |
| R6 | 709 |
| Validation | 710 |
| R1 | 710 |
| RSVP-TE Dynamic Facility Backup LSP (RSVP Auto Bypass) | 721 |
| Overview | 721 |
| Feature Characteristics | 721 |
| Benefits | 721 |
| Prerequisites | 721 |
| Define Interfaces and Loopback Addresses | 721 |
| Configure IGP for Dynamic Routing | 722 |

| | |
|---|-----|
| Configure RSVP for Efficient Network Operation | 722 |
| Configuration for RSVP Auto Bypass | 723 |
| Topology | 723 |
| Configuring RSVP Auto Bypass | 723 |
| Auto Bypass and Manual Bypass Co-exists with Perfect Protection | 733 |
| Auto Bypass with Node Protection and Manual Bypass with Link Protection | 735 |
| RSVP Auto Bypass Commands | 737 |
| auto-bypass | 738 |
| auto-bypass enable | 739 |
| inactivity-timer | 740 |
| attributes | 741 |
| protection-capability | 742 |
| bandwidth | 743 |
| setup-priority | 744 |
| hold-priority | 745 |
| hop-limit | 746 |
| reoptimize | 747 |
| exclude-address | 748 |
| show rsvp router-id-table | 750 |
| show rsvp auto-bypass-group | 751 |
| RSVP Graceful Restart Configuration | 752 |
| Overview | 752 |
| Topology | 752 |
| Configuration | 752 |
| PE1-NSM | 752 |
| PE1-RSVP-TE | 752 |
| PE1-OSPF | 753 |
| P - NSM | 753 |
| P - RSVP-TE | 753 |
| P - OSPF | 754 |
| PE3 - NSM | 754 |
| PE3- RSVP-TE | 754 |
| PE3 - OSPF | 755 |
| PE1 - RSVP-Path | 755 |
| PE3 - RSVP-Path | 755 |
| PE1 – RSVP-GR | 756 |
| P1 – RSVP-GR | 756 |
| PE3 – RSVP-GR | 756 |
| Validation | 756 |
| Fast Reroute Configuration (One-to-one Method) | 758 |
| RTR-29 Configuration | 758 |
| RTR-43 Configuration | 761 |
| RTR-42 Configuration | 762 |
| RTR-41 Configuration | 764 |
| RTR-31 Configuration | 765 |
| RTR-30 Configuration | 766 |
| Validation | 768 |

| | |
|--|-----|
| RSVP Detour Over Ring Topology | 771 |
| Overview | 771 |
| Prerequisite | 772 |
| Configuration | 772 |
| Implementation Examples | 788 |
| RSVP Detour Over Ring Topology Command | 788 |
| Abbreviations | 788 |
| Glossary | 789 |
| Auto-Bandwidth with RSVP-TE | 790 |
| Overview | 790 |
| Feature Characteristics | 790 |
| Benefits | 791 |
| Prerequisites | 791 |
| Define Interfaces and Loopback Addresses | 791 |
| Configure IGP for Dynamic Routing | 792 |
| Configure RSVP for Efficient Network Operation | 792 |
| Configure the RSVP Primary Path and Trunk | 793 |
| Configuration for RSVP Auto-Bandwidth | 793 |
| Topology | 793 |
| PE1 | 793 |
| Running Configuration on PE1 Router | 794 |
| Validation | 794 |
| Boot on PE1 | 800 |
| Commands for RSVP Auto-Bandwidth | 801 |
| rsvp-auto-bandwidth | 802 |
| sample-interval | 803 |
| adjust-interval | 804 |
| minimum-bandwidth | 805 |
| maximum-bandwidth | 806 |
| initial-bandwidth | 807 |
| underflow-threshold | 808 |
| overflow-threshold | 810 |
| underflow-threshold-activate-bandwidth | 812 |
| overflow-threshold-activate-bandwidth | 813 |
| underflow-limit | 814 |
| overflow-limit | 815 |
| maximum-bandwidth-exceed-limit | 816 |
| maximum-bandwidth-exceed-action | 817 |
| resignal-failure-action | 818 |
| sync-bandwidth | 819 |
| monitor-bandwidth | 820 |
| minimum-samples | 821 |
| auto-bandwidth | 823 |
| auto-bandwidth-on-boot | 824 |
| force-auto-bandwidth-adjustment | 825 |
| clear rsvp auto-bandwidth | 826 |
| clear rsvp trunk auto-bandwidth-statistics | 827 |

| | |
|---|------------|
| show rsvp auto-bandwidth | 828 |
| show rsvp auto-bandwidth detail | 829 |
| show rsvp auto-bypass-group | 830 |
| show rsvp trunk auto-bandwidth | 831 |
| show rsvp trunk auto-bandwidth detail | 832 |
| RSVP LSP Reoptimization | 834 |
| Overview | 834 |
| Characteristics | 834 |
| Benefits | 834 |
| Limitations | 834 |
| Prerequisites | 834 |
| Configuration | 835 |
| Topology | 835 |
| P1 | 835 |
| P2 | 836 |
| PE1 | 837 |
| PE2 | 838 |
| Running Configuration | 840 |
| P1 | 840 |
| P2 | 841 |
| PE1 | 841 |
| PE2 | 843 |
| Validation | 844 |
| Base Configuration Validations | 844 |
| Scenario 1: Verify Reoptimization for Primary Tunnels | 846 |
| Scenario 2: Verify Re-optimization for Secondary Tunnels | 850 |
| Scenario 3: Verify Reoptimization for Detour Tunnels | 854 |
| Scenario 4: Verify Reoptimization for Auto-bypass Tunnels | 857 |
| Implementation Examples | 861 |
| Revised Commands | 861 |
| Troubleshooting | 861 |
| Glossary | 862 |
| RSVP-TE Command Reference | 863 |
| RSVP-TE Commands | 869 |
| A.B.C.D | 872 |
| Command Syntax | 872 |
| Parameters | 872 |
| Default | 872 |
| Command Mode | 872 |
| Applicability | 872 |
| Examples | 872 |
| ack-wait-timeout | 873 |
| Command Syntax | 873 |
| Parameters | 873 |
| Default | 873 |
| Command Mode | 873 |

| | |
|---------------------------|-----|
| Applicability | 873 |
| Examples | 873 |
| clear rsvp session | 874 |
| Command Syntax | 874 |
| Parameters | 874 |
| Command Mode | 874 |
| Applicability | 874 |
| Examples | 874 |
| clear rsvp neighbor | 875 |
| Command Syntax | 875 |
| Parameters | 875 |
| Command Mode | 875 |
| Applicability | 875 |
| Examples | 875 |
| clear rsvp trunk | 876 |
| Command Syntax | 876 |
| Parameters | 876 |
| Command Mode | 876 |
| Applicability | 876 |
| Examples | 877 |
| cspf | 878 |
| Command Syntax | 878 |
| Parameters | 878 |
| Default | 878 |
| Command Mode | 878 |
| Applicability | 878 |
| Example | 878 |
| debug rsvp all | 879 |
| Command Syntax | 879 |
| Parameters | 879 |
| Command Mode | 879 |
| Applicability | 879 |
| Examples | 879 |
| debug rsvp cspf | 880 |
| Command Syntax | 880 |
| Parameters | 880 |
| Command Mode | 880 |
| Applicability | 880 |
| Examples | 880 |
| debug rsvp events | 881 |
| Command Syntax | 881 |
| Parameters | 881 |
| Command Mode | 881 |
| Applicability | 881 |
| Examples | 881 |
| debug rsvp fsm | 882 |
| Command Syntax | 882 |

| | |
|---------------------------------|-----|
| Parameters | 882 |
| Command Mode | 882 |
| Applicability | 882 |
| Examples | 882 |
| debug rsvp hexdump | 883 |
| Command Syntax | 883 |
| Parameters | 883 |
| Command Mode | 883 |
| Applicability | 883 |
| Examples | 883 |
| debug rsvp nsm | 884 |
| Command Syntax | 884 |
| Parameters | 884 |
| Command Mode | 884 |
| Applicability | 884 |
| Examples | 884 |
| debug rsvp packet | 885 |
| Command Syntax | 885 |
| Parameters | 885 |
| Command Mode | 885 |
| Applicability | 885 |
| Examples | 885 |
| description (rsvp-bypass) | 886 |
| Command Syntax | 886 |
| Parameters | 886 |
| Default | 886 |
| Command Mode | 886 |
| Applicability | 886 |
| Examples | 886 |
| description (rsvp-path) | 887 |
| Command Syntax | 887 |
| Parameters | 887 |
| Default | 887 |
| Command Mode | 887 |
| Applicability | 887 |
| Examples | 887 |
| description (rsvp-trunk) | 888 |
| Command Syntax | 888 |
| Parameters | 888 |
| Default | 888 |
| Command Mode | 888 |
| Applicability | 888 |
| Examples | 888 |
| disable-rsvp | 889 |
| Command Syntax | 889 |
| Parameters | 889 |
| Default | 889 |

| | |
|--------------------------------------|-----|
| Command Mode | 889 |
| Applicability | 889 |
| Examples | 889 |
| enable-rsvp | 890 |
| Command Syntax | 890 |
| Parameters | 890 |
| Default | 890 |
| Command Mode | 890 |
| Applicability | 890 |
| Examples | 890 |
| entropy-label-capability | 891 |
| Command Syntax | 891 |
| Parameters | 891 |
| Default | 891 |
| Command Mode | 891 |
| Applicability | 891 |
| Examples | 891 |
| explicit-null | 892 |
| Command Syntax | 892 |
| Parameters | 892 |
| Default | 892 |
| Command Mode | 892 |
| Applicability | 892 |
| Examples | 892 |
| ext-tunnel-id A.B.C.D | 893 |
| Command Syntax | 893 |
| Parameters | 893 |
| Default | 893 |
| Command Mode | 893 |
| Applicability | 893 |
| Examples | 893 |
| from A.B.C.D | 894 |
| Command Syntax | 894 |
| Parameters | 894 |
| Default | 894 |
| Command Mode | 894 |
| Applicability | 894 |
| Examples | 894 |
| graceful-restart | 895 |
| Command Syntax | 895 |
| Parameters | 895 |
| Default | 895 |
| Command Mode | 895 |
| Applicability | 895 |
| Examples | 895 |
| graceful-restart recovery-time | 896 |
| Command Syntax | 896 |

| | |
|-------------------------------------|-----|
| Parameters | 896 |
| Default | 896 |
| Command Mode | 896 |
| Applicability | 896 |
| Examples | 896 |
| graceful-restart restart-time | 897 |
| Command Syntax | 897 |
| Parameters | 897 |
| Default | 897 |
| Command Mode | 897 |
| Applicability | 897 |
| Examples | 897 |
| hello-interval | 898 |
| Command Syntax | 898 |
| Parameter | 898 |
| Default | 898 |
| Command Mode | 898 |
| Applicability | 898 |
| Examples | 898 |
| hello-receipt | 899 |
| Command Syntax | 899 |
| Parameters | 899 |
| Default | 899 |
| Command Mode | 899 |
| Applicability | 899 |
| Examples | 899 |
| hello-timeout | 900 |
| Command Syntax | 900 |
| Parameter | 900 |
| Default | 900 |
| Command Mode | 900 |
| Applicability | 900 |
| Examples | 900 |
| keep-multiplier | 901 |
| Command Syntax | 901 |
| Parameters | 901 |
| Default | 901 |
| Command Mode | 901 |
| Applicability | 901 |
| Examples | 901 |
| loop-detection | 902 |
| Command Syntax | 902 |
| Parameters | 902 |
| Default | 902 |
| Command Mode | 902 |
| Applicability | 902 |
| Examples | 902 |

| | |
|--------------------------------|-----|
| lsp-metric | 903 |
| Command Syntax | 903 |
| Parameters | 903 |
| Command Mode | 903 |
| Examples | 903 |
| lsp-reoptimization-timer | 904 |
| Command Syntax | 904 |
| Parameter | 904 |
| Default | 904 |
| Command Mode | 904 |
| Applicability | 904 |
| Example | 904 |
| map-route A.B.C.D | 905 |
| Command Syntax | 905 |
| Parameters | 905 |
| Default | 905 |
| Command Mode | 905 |
| Applicability | 905 |
| Example | 905 |
| neighbor A.B.C.D | 906 |
| Command Syntax | 906 |
| Parameters | 906 |
| Default | 906 |
| Command Mode | 906 |
| Applicability | 906 |
| Examples | 906 |
| neighbor X:X::X:X | 907 |
| Command Syntax | 907 |
| Parameters | 907 |
| Default | 907 |
| Command Mode | 907 |
| Applicability | 907 |
| Examples | 907 |
| no cspf | 908 |
| Command Syntax | 908 |
| Parameters | 908 |
| Default | 908 |
| Command Mode | 908 |
| Applicability | 908 |
| Example | 908 |
| no loop-detection | 909 |
| Command Syntax | 909 |
| Parameters | 909 |
| Default | 909 |
| Command Mode | 909 |
| Applicability | 909 |
| Examples | 909 |

| | |
|-------------------------------|-----|
| no php | 910 |
| Command Syntax | 910 |
| Parameters | 910 |
| Default | 910 |
| Command Mode | 910 |
| Applicability | 910 |
| Examples | 910 |
| no primary | 911 |
| Command Syntax | 911 |
| Parameters | 911 |
| Default | 911 |
| Command Mode | 911 |
| Applicability | 911 |
| Example | 911 |
| no primary affinity | 912 |
| Command Syntax | 912 |
| Parameters | 912 |
| Default | 912 |
| Command Mode | 912 |
| Applicability | 912 |
| Example | 912 |
| no primary cspf | 913 |
| Command Syntax | 913 |
| Parameters | 913 |
| Default | 913 |
| Command Mode | 913 |
| Applicability | 913 |
| Example | 913 |
| no primary record | 914 |
| Command Syntax | 914 |
| Parameters | 914 |
| Default | 914 |
| Command Mode | 914 |
| Applicability | 914 |
| Examples | 914 |
| no record | 915 |
| Command Syntax | 915 |
| Parameters | 915 |
| Default | 915 |
| Command Mode | 915 |
| Examples | 915 |
| no refresh-path-parsing | 916 |
| Command Syntax | 916 |
| Parameters | 916 |
| Default | 916 |
| Command Mode | 916 |
| Applicability | 916 |

| | |
|--------------------------------|-----|
| Example | 916 |
| no refresh-resv-parsing | 917 |
| Command Syntax | 917 |
| Parameters | 917 |
| Default | 917 |
| Command Mode | 917 |
| Applicability | 917 |
| Example | 917 |
| path-option dynamic pce | 918 |
| Command Syntax | 918 |
| Parameters | 918 |
| Command Mode | 918 |
| Applicability | 918 |
| Example | 918 |
| php | 919 |
| Command Syntax | 919 |
| Parameters | 919 |
| Default | 919 |
| Command Mode | 919 |
| Applicability | 919 |
| Examples | 919 |
| primary ADMIN-GROUP-NAME | 920 |
| Command Syntax | 920 |
| Parameters | 920 |
| Default | 920 |
| Command Mode | 920 |
| Applicability | 920 |
| Examples | 920 |
| primary affinity | 922 |
| Command Syntax | 922 |
| Parameters | 922 |
| Default | 922 |
| Command Mode | 922 |
| Applicability | 922 |
| Example | 922 |
| primary bandwidth | 923 |
| Command Syntax | 923 |
| Parameter | 923 |
| Default | 923 |
| Command Mode | 923 |
| Applicability | 923 |
| Examples | 923 |
| primary cspf | 924 |
| Command Syntax | 924 |
| Parameters | 924 |
| Default | 924 |
| Command Mode | 924 |

| | |
|--------------------------------|-----|
| Applicability | 924 |
| Example | 924 |
| primary cspf-retry-limit | 925 |
| Command Syntax | 925 |
| Parameter | 925 |
| Default | 925 |
| Command Mode | 925 |
| Applicability | 925 |
| Examples | 925 |
| primary cspf-retry-timer | 926 |
| Command Syntax | 926 |
| Parameter | 926 |
| Default | 926 |
| Command Mode | 926 |
| Applicability | 926 |
| Examples | 926 |
| primary filter | 927 |
| Command Syntax | 927 |
| Parameters | 927 |
| Default | 927 |
| Command Mode | 927 |
| Applicability | 927 |
| Examples | 927 |
| primary hold-priority | 928 |
| Command Syntax | 928 |
| Parameter | 928 |
| Default | 928 |
| Command Mode | 928 |
| Applicability | 928 |
| Examples | 928 |
| primary hop-limit | 929 |
| Command Syntax | 929 |
| Parameters | 929 |
| Default | 929 |
| Command Mode | 929 |
| Applicability | 929 |
| Examples | 929 |
| primary label-record | 930 |
| Command Syntax | 930 |
| Parameters | 930 |
| Default | 930 |
| Command Mode | 930 |
| Applicability | 930 |
| Examples | 930 |
| primary path | 931 |
| Command Syntax | 931 |
| Parameters | 931 |

| | |
|----------------------------------|-----|
| Default | 931 |
| Command Mode | 931 |
| Applicability | 931 |
| Examples | 931 |
| primary record | 932 |
| Command Syntax | 932 |
| Parameters | 932 |
| Default | 932 |
| Command Mode | 932 |
| Applicability | 932 |
| Examples | 932 |
| primary retry-limit | 933 |
| Command Syntax | 933 |
| Parameter | 933 |
| Default | 933 |
| Command Mode | 933 |
| Applicability | 933 |
| Examples | 933 |
| primary retry-timer | 934 |
| Command Syntax | 934 |
| Parameter | 934 |
| Default | 934 |
| Command Mode | 934 |
| Applicability | 934 |
| Examples | 934 |
| primary reuse-route-record | 935 |
| Command Syntax | 935 |
| Parameters | 935 |
| Default | 935 |
| Command Mode | 935 |
| Applicability | 935 |
| Examples | 935 |
| primary setup-priority | 936 |
| Command Syntax | 936 |
| Parameters | 936 |
| Default | 936 |
| Command Mode | 936 |
| Applicability | 936 |
| Examples | 936 |
| primary traffic | 937 |
| Command Syntax | 937 |
| Parameters | 937 |
| Default | 937 |
| Command Mode | 937 |
| Applicability | 937 |
| Examples | 937 |
| refresh-time | 938 |

| | |
|-----------------------------|-----|
| Command Syntax | 938 |
| Parameter | 938 |
| Default | 938 |
| Command Mode | 938 |
| Applicability | 938 |
| Examples | 938 |
| refresh-path-parsing | 939 |
| Command Syntax | 939 |
| Parameters | 939 |
| Default | 939 |
| Command Mode | 939 |
| Applicability | 939 |
| Example | 939 |
| refresh-resv-parsing | 940 |
| Command Syntax | 940 |
| Parameters | 940 |
| Default | 940 |
| Command Mode | 940 |
| Applicability | 940 |
| Example | 940 |
| reoptimize | 941 |
| Command Syntax | 941 |
| Parameters | 941 |
| Command Mode | 941 |
| Applicability | 941 |
| Examples | 941 |
| restart rsvp graceful | 942 |
| Command Syntax | 942 |
| Parameter | 942 |
| Command Mode | 942 |
| Applicability | 942 |
| Example | 942 |
| revert-timer | 943 |
| Command Syntax | 943 |
| Parameters | 943 |
| Default | 943 |
| Command Mode | 943 |
| Applicability | 943 |
| Examples | 943 |
| router rsvp | 944 |
| Command Syntax | 944 |
| Parameters | 944 |
| Default | 944 |
| Command Mode | 944 |
| Applicability | 944 |
| Example | 944 |
| rsvp hello-interval | 945 |

| | |
|----------------------------|-----|
| Command Syntax | 945 |
| Parameter | 945 |
| Default | 945 |
| Command Mode | 945 |
| Applicability | 945 |
| Example | 945 |
| rsvp hello-receipt | 946 |
| Command Syntax | 946 |
| Parameters | 946 |
| Default | 946 |
| Command Mode | 946 |
| Applicability | 946 |
| Example | 946 |
| rsvp hello-timeout | 947 |
| Command Syntax | 947 |
| Parameters | 947 |
| Default | 947 |
| Command Mode | 947 |
| Applicability | 947 |
| Examples | 947 |
| rsvp keep-multiplier | 948 |
| Command Syntax | 948 |
| Parameters | 948 |
| Default | 948 |
| Command Mode | 948 |
| Applicability | 948 |
| Examples | 948 |
| rsvp refresh-time | 949 |
| Command Syntax | 949 |
| Parameter | 949 |
| Default | 949 |
| Command Mode | 949 |
| Applicability | 949 |
| Examples | 949 |
| rsvp-path | 950 |
| Command Syntax | 950 |
| Parameter | 950 |
| Default | 950 |
| Command Mode | 950 |
| Applicability | 950 |
| Example | 950 |
| rsvp-trunk | 951 |
| Command Syntax | 951 |
| Parameters | 951 |
| Default | 951 |
| Command Mode | 951 |
| Applicability | 951 |

| | |
|---|-----|
| Examples | 951 |
| rsvp-session force-reoptimize | 952 |
| Command Syntax | 952 |
| Parameters | 952 |
| Default | 952 |
| Command Mode | 952 |
| Applicability | 952 |
| Examples | 952 |
| rsvp-trunk force-switchover-secondary | 953 |
| Command Syntax | 953 |
| Parameters | 953 |
| Default | 953 |
| Command Mode | 953 |
| Applicability | 953 |
| Examples | 953 |
| rsvp-trunk-restart | 954 |
| Command Syntax | 954 |
| Parameters | 954 |
| Default | 954 |
| Command Mode | 954 |
| Applicability | 954 |
| Examples | 954 |
| secondary | 955 |
| Command Syntax | 955 |
| Parameters | 955 |
| Default | 955 |
| Command Mode | 955 |
| Applicability | 955 |
| Example | 955 |
| secondary ADMIN-GROUP-NAME | 956 |
| Command Syntax | 956 |
| Parameters | 956 |
| Default | 956 |
| Command Mode | 956 |
| Applicability | 956 |
| Examples | 956 |
| secondary bandwidth | 958 |
| Command Syntax | 958 |
| Parameter | 958 |
| Default | 958 |
| Command Mode | 958 |
| Applicability | 958 |
| Examples | 958 |
| secondary cspf | 959 |
| Command Syntax | 959 |
| Parameters | 959 |
| Default | 959 |

| | |
|----------------------------------|-----|
| Command Mode | 959 |
| Applicability | 959 |
| Example | 959 |
| secondary cspf-retry-limit | 960 |
| Command Syntax | 960 |
| Parameter | 960 |
| Default | 960 |
| Command Mode | 960 |
| Applicability | 960 |
| Examples | 960 |
| secondary cspf-retry-timer | 961 |
| Command Syntax | 961 |
| Parameters | 961 |
| Default | 961 |
| Command Mode | 961 |
| Applicability | 961 |
| Examples | 961 |
| secondary filter | 962 |
| Command Syntax | 962 |
| Parameters | 962 |
| Default | 962 |
| Command Mode | 962 |
| Applicability | 962 |
| Usage | 962 |
| Examples | 962 |
| secondary hold-priority | 963 |
| Command Syntax | 963 |
| Parameter | 963 |
| Default | 963 |
| Command Mode | 963 |
| Applicability | 963 |
| Examples | 963 |
| secondary hop-limit | 964 |
| Command Syntax | 964 |
| Parameter | 964 |
| Default | 964 |
| Command Mode | 964 |
| Applicability | 964 |
| Examples | 964 |
| secondary label-record | 965 |
| Command Syntax | 965 |
| Default | 965 |
| Command Mode | 965 |
| Applicability | 965 |
| Examples | 965 |
| no secondary affinity | 966 |
| Command Syntax | 966 |

| | |
|---|-----|
| Parameters | 966 |
| Default | 966 |
| Command Mode | 966 |
| Applicability | 966 |
| Example | 966 |
| no secondary cspf | 967 |
| Command Syntax | 967 |
| Parameters | 967 |
| Default | 967 |
| Command Mode | 967 |
| Applicability | 967 |
| Example | 967 |
| no secondary record | 968 |
| Command Syntax | 968 |
| Parameters | 968 |
| Default | 968 |
| Command Mode | 968 |
| Applicability | 968 |
| Examples | 968 |
| secondary path | 969 |
| Command Syntax | 969 |
| Parameters | 969 |
| Default | 969 |
| Command Mode | 969 |
| Applicability | 969 |
| Examples | 969 |
| secondary-priority path | 970 |
| Command Syntax | 970 |
| Parameters | 970 |
| Default | 970 |
| Command Mode | 970 |
| Applicability | 970 |
| Examples | 970 |
| secondary-priority hold-priority | 971 |
| Command Syntax | 971 |
| Parameters | 971 |
| Default | 971 |
| Command Mode | 971 |
| Applicability | 971 |
| Examples | 971 |
| secondary-priority setup-priority | 972 |
| Command Syntax | 972 |
| Parameters | 972 |
| Default | 972 |
| Command Mode | 972 |
| Applicability | 972 |
| Examples | 972 |

| | |
|---------------------------------------|-----|
| secondary-priority label-record | 973 |
| Command Syntax | 973 |
| Parameters | 973 |
| Default | 973 |
| Command Mode | 973 |
| Applicability | 973 |
| Examples | 973 |
| secondary-priority hop-limit | 974 |
| Command Syntax | 974 |
| Parameters | 974 |
| Default | 974 |
| Command Mode | 974 |
| Applicability | 974 |
| Examples | 974 |
| secondary-priority bandwidth | 975 |
| Command Syntax | 975 |
| Parameters | 975 |
| Default | 975 |
| Command Mode | 975 |
| Applicability | 975 |
| Examples | 975 |
| secondary record | 976 |
| Command Syntax | 976 |
| Parameters | 976 |
| Default | 976 |
| Command Mode | 976 |
| Applicability | 976 |
| Examples | 976 |
| secondary retry-limit | 977 |
| Command Syntax | 977 |
| Parameter | 977 |
| Default | 977 |
| Command Mode | 977 |
| Applicability | 977 |
| Examples | 977 |
| secondary retry-timer | 978 |
| Command Syntax | 978 |
| Parameter | 978 |
| Default | 978 |
| Command Mode | 978 |
| Applicability | 978 |
| Examples | 978 |
| secondary reuse-route-record | 979 |
| Command Syntax | 979 |
| Parameters | 979 |
| Default | 979 |
| Command Mode | 979 |

| | |
|--------------------------------|-----|
| Applicability | 979 |
| Examples | 979 |
| secondary setup-priority | 980 |
| Command Syntax | 980 |
| Parameters | 980 |
| Default | 980 |
| Command Mode | 980 |
| Applicability | 980 |
| Examples | 980 |
| secondary traffic | 981 |
| Command Syntax | 981 |
| Parameters | 981 |
| Default | 981 |
| Command Mode | 981 |
| Applicability | 981 |
| Examples | 981 |
| session state timeout | 982 |
| Command Syntax | 982 |
| Parameters | 982 |
| Default | 982 |
| Command Mode | 982 |
| Applicability | 982 |
| Example | 982 |
| shutdown | 983 |
| Command Syntax | 983 |
| Parameters | 983 |
| Default | 983 |
| Command Mode | 983 |
| Applicability | 983 |
| Examples | 983 |
| snmp restart rsvp | 984 |
| Command Syntax | 984 |
| Parameters | 984 |
| Default | 984 |
| Command Mode | 984 |
| Applicability | 984 |
| Examples | 984 |
| to A.B.C.D | 985 |
| Command Syntax | 985 |
| Parameters | 985 |
| Default | 985 |
| Command Mode | 985 |
| Applicability | 985 |
| Examples | 985 |
| update-type | 986 |
| Command Syntax | 986 |
| Parameters | 986 |

| | |
|---|-----|
| Default | 986 |
| Command Mode | 986 |
| Applicability | 986 |
| Examples | 986 |
| Fast Reroute Commands | 987 |
| detour-allow-primary-upstream-path | 988 |
| Command Syntax | 988 |
| Parameters | 988 |
| Default | 988 |
| Command Mode | 988 |
| Applicability | 988 |
| Examples | 988 |
| default-frr-protection | 989 |
| Command Syntax | 989 |
| Parameters | 989 |
| Default | 989 |
| Command Mode | 989 |
| Applicability | 989 |
| Examples | 989 |
| detour-identification | 990 |
| Command Syntax | 990 |
| Parameters | 990 |
| Default | 990 |
| Command Mode | 990 |
| Applicability | 990 |
| Examples | 990 |
| primary fast-reroute admin-group-name | 991 |
| Command Syntax | 991 |
| Parameter | 991 |
| Default | 991 |
| Command Mode | 991 |
| Applicability | 991 |
| Examples | 991 |
| primary fast-reroute bandwidth | 992 |
| Command Syntax | 992 |
| Parameter | 992 |
| Default | 992 |
| Command Mode | 992 |
| Applicability | 992 |
| Examples | 992 |
| primary fast-reroute hold-priority | 993 |
| Command Syntax | 993 |
| Parameter | 993 |
| Default | 993 |
| Command Mode | 993 |
| Applicability | 993 |
| Examples | 993 |

| | |
|--|------|
| primary fast-reroute hop-limit | 994 |
| Command Syntax | 994 |
| Parameter | 994 |
| Default | 994 |
| Command Mode | 994 |
| Applicability | 994 |
| Examples | 994 |
| primary fast-reroute node-protection | 995 |
| Command Syntax | 995 |
| Parameters | 995 |
| Default | 995 |
| Command Mode | 995 |
| Applicability | 995 |
| Examples | 995 |
| primary fast-reroute protection | 996 |
| Parameters | 996 |
| Command Syntax | 996 |
| Parameters | 996 |
| Default | 996 |
| Command Mode | 996 |
| Applicability | 996 |
| Examples | 996 |
| primary fast-reroute setup-priority | 997 |
| Command Syntax | 997 |
| Parameter | 997 |
| Default | 997 |
| Command Mode | 997 |
| Applicability | 997 |
| Examples | 997 |
| Refresh Reduction Commands | 998 |
| ack-send-timer | 999 |
| Command Syntax | 999 |
| Parameter | 999 |
| Default | 999 |
| Command Mode | 999 |
| Applicability | 999 |
| Examples | 999 |
| ack-wait-timeout | 1000 |
| Command Syntax | 1000 |
| Parameters | 1000 |
| Default | 1000 |
| Command Mode | 1000 |
| Applicability | 1000 |
| Examples | 1000 |
| message-ack | 1001 |
| Command Syntax | 1001 |
| Parameters | 1001 |

| | |
|--------------------------------|------|
| Default | 1001 |
| Command Mode | 1001 |
| Applicability | 1001 |
| Examples | 1001 |
| refresh-reduction | 1002 |
| Command Syntax | 1002 |
| Parameters | 1002 |
| Default | 1002 |
| Command Mode | 1002 |
| Applicability | 1002 |
| Examples | 1002 |
| rsvp ack-wait-timeout | 1003 |
| Command Syntax | 1003 |
| Parameters | 1003 |
| Default | 1003 |
| Command Mode | 1003 |
| Applicability | 1003 |
| Examples | 1003 |
| rsvp message-ack | 1004 |
| Command Syntax | 1004 |
| Parameters | 1004 |
| Default | 1004 |
| Command Mode | 1004 |
| Applicability | 1004 |
| Examples | 1004 |
| rsvp refresh-reduction | 1005 |
| Command Syntax | 1005 |
| Parameters | 1005 |
| Default | 1005 |
| Command Mode | 1005 |
| Applicability | 1005 |
| Examples | 1005 |
| Facility Backup Commands | 1006 |
| backup-bw-type | 1007 |
| Command Syntax | 1007 |
| Parameters | 1007 |
| Default | 1007 |
| Command Mode | 1007 |
| Applicability | 1007 |
| Examples | 1007 |
| bandwidth | 1008 |
| Command Syntax | 1008 |
| Parameter | 1008 |
| Default | 1008 |
| Command Mode | 1008 |
| Applicability | 1008 |
| Examples | 1008 |

| | |
|--------------------------------------|------|
| bypass-lsp-addr-query-interval | 1009 |
| Command Syntax | 1009 |
| Parameter | 1009 |
| Default | 1009 |
| Command Mode | 1009 |
| Applicability | 1009 |
| Examples | 1009 |
| cspf-retry-limit | 1010 |
| Command Syntax | 1010 |
| Parameter | 1010 |
| Default | 1010 |
| Command Mode | 1010 |
| Applicability | 1010 |
| Examples | 1010 |
| cspf-retry-timer | 1011 |
| Command Syntax | 1011 |
| Parameter | 1011 |
| Default | 1011 |
| Command Mode | 1011 |
| Applicability | 1011 |
| Examples | 1011 |
| filter | 1012 |
| Command Syntax | 1012 |
| Parameters | 1012 |
| Default | 1012 |
| Command Mode | 1012 |
| Applicability | 1012 |
| Examples | 1012 |
| hold-priority | 1013 |
| Command Syntax | 1013 |
| Parameters | 1013 |
| Default | 1013 |
| Command Mode | 1013 |
| Applicability | 1013 |
| Examples | 1013 |
| hop-limit | 1014 |
| Command Syntax | 1014 |
| Parameters | 1014 |
| Default | 1014 |
| Command Mode | 1014 |
| Applicability | 1014 |
| Examples | 1014 |
| label-record | 1015 |
| Command Syntax | 1015 |
| Parameters | 1015 |
| Default | 1015 |
| Command Mode | 1015 |

| | |
|--------------------------|------|
| Applicability | 1015 |
| Examples | 1015 |
| no record | 1016 |
| Command Syntax | 1016 |
| Parameters | 1016 |
| Default | 1016 |
| Command Mode | 1016 |
| Examples | 1016 |
| path | 1017 |
| Command Syntax | 1017 |
| Parameters | 1017 |
| Default | 1017 |
| Command Mode | 1017 |
| Applicability | 1017 |
| Examples | 1017 |
| preemption-type | 1018 |
| Command Syntax | 1018 |
| Parameters | 1018 |
| Default | 1018 |
| Command Mode | 1018 |
| Applicability | 1018 |
| Examples | 1018 |
| record | 1019 |
| Command Syntax | 1019 |
| Parameters | 1019 |
| Default | 1019 |
| Command Mode | 1019 |
| Applicability | 1019 |
| Examples | 1019 |
| retry-limit | 1020 |
| Command Syntax | 1020 |
| Parameter | 1020 |
| Default | 1020 |
| Command Mode | 1020 |
| Applicability | 1020 |
| Examples | 1020 |
| retry-timer | 1021 |
| Command Syntax | 1021 |
| Parameters | 1021 |
| Default | 1021 |
| Command Mode | 1021 |
| Applicability | 1021 |
| Examples | 1021 |
| reuse-route-record | 1022 |
| Command Syntax | 1022 |
| Parameters | 1022 |
| Default | 1022 |

| | |
|--|------|
| Command Mode | 1022 |
| Applicability | 1022 |
| Examples | 1022 |
| rsvp-bypass | 1023 |
| Command Syntax | 1023 |
| Parameters | 1023 |
| Default | 1023 |
| Command Mode | 1023 |
| Applicability | 1023 |
| Examples | 1023 |
| setup-priority | 1024 |
| Command Syntax | 1024 |
| Parameters | 1024 |
| Default | 1024 |
| Command Mode | 1024 |
| Applicability | 1024 |
| Examples | 1024 |
| shutdown | 1025 |
| Command Syntax | 1025 |
| Parameters | 1025 |
| Default | 1025 |
| Command Mode | 1025 |
| Applicability | 1025 |
| Examples | 1025 |
| to A.B.C.D | 1026 |
| Command Syntax | 1026 |
| Parameters | 1026 |
| Default | 1026 |
| Command Mode | 1026 |
| Applicability | 1026 |
| Examples | 1026 |
| traffic | 1027 |
| Command Syntax | 1027 |
| Parameters | 1027 |
| Default | 1027 |
| Command Mode | 1027 |
| Applicability | 1027 |
| Examples | 1027 |
| Differentiated Services Commands | 1028 |
| map-route A.B.C.D | 1029 |
| Command Syntax | 1029 |
| Parameters | 1029 |
| Default | 1029 |
| Command Mode | 1029 |
| Applicability | 1029 |
| Example | 1029 |
| map-route X:X::X:X | 1030 |

| | |
|--------------------------------|------|
| Command Syntax | 1030 |
| Parameters | 1030 |
| Default | 1030 |
| Command Mode | 1030 |
| Applicability | 1030 |
| Example | 1030 |
| override-diffserv | 1031 |
| Command Syntax | 1031 |
| Parameters | 1031 |
| Default | 1031 |
| Command Mode | 1031 |
| Applicability | 1031 |
| Example | 1031 |
| primary class-to-exp-bit | 1032 |
| Command Syntax | 1032 |
| Parameters | 1032 |
| Default | 1032 |
| Command Mode | 1032 |
| Applicability | 1032 |
| Examples | 1032 |
| primary elsp-signaled | 1033 |
| Command Syntax | 1033 |
| Parameter | 1033 |
| Default | 1033 |
| Command Mode | 1033 |
| Applicability | 1033 |
| Examples | 1033 |
| primary llsp | 1034 |
| Command Syntax | 1034 |
| Parameters | 1034 |
| Default | 1034 |
| Command Mode | 1034 |
| Applicability | 1034 |
| Examples | 1034 |
| secondary elsp-signaled | 1035 |
| Command Syntax | 1035 |
| Parameters | 1035 |
| Default | 1035 |
| Command Mode | 1035 |
| Applicability | 1035 |
| Examples | 1035 |
| secondary llsp | 1036 |
| Command Syntax | 1036 |
| Parameters | 1036 |
| Default | 1036 |
| Command Mode | 1036 |
| Applicability | 1036 |

| | |
|--|------|
| Example | 1036 |
| secondary map class | 1037 |
| Command Syntax | 1037 |
| Parameters | 1037 |
| Default | 1037 |
| Command Mode | 1037 |
| Applicability | 1037 |
| Examples | 1037 |
| show rsvp diffserv-info | 1038 |
| Command Syntax | 1038 |
| Parameters | 1038 |
| Command Mode | 1038 |
| Applicability | 1038 |
| Example | 1038 |
| RSVP Multipath Commands | 1039 |
| description | 1040 |
| Command Syntax | 1040 |
| Parameter | 1040 |
| Default | 1040 |
| Command Mode | 1040 |
| Applicability | 1040 |
| Examples | 1040 |
| multipath-group | 1041 |
| Command Syntax | 1041 |
| Parameter | 1041 |
| Default | 1041 |
| Command Mode | 1041 |
| Applicability | 1041 |
| Examples | 1041 |
| rsvp-multipath | 1042 |
| Command Syntax | 1042 |
| Parameter | 1042 |
| Default | 1042 |
| Command Mode | 1042 |
| Applicability | 1042 |
| Examples | 1042 |
| to A.B.C.D | 1043 |
| Command Syntax | 1043 |
| Parameter | 1043 |
| Default | 1043 |
| Command Mode | 1043 |
| Applicability | 1043 |
| Examples | 1043 |
| show mpls counters rsvp multipath-name | 1044 |
| Command Syntax | 1044 |
| Parameter | 1044 |
| Command Mode | 1044 |

| | |
|--|------|
| Applicability | 1044 |
| Examples | 1044 |
| show mpls rsvp-multipath | 1045 |
| Command Syntax | 1045 |
| Parameter | 1045 |
| Command Mode | 1045 |
| Applicability | 1045 |
| Examples | 1045 |
| show rsvp multipath | 1046 |
| Command Syntax | 1046 |
| Parameter | 1046 |
| Command Mode | 1046 |
| Applicability | 1046 |
| Examples | 1046 |
| show running-config rsvp-multipath | 1047 |
| Command Syntax | 1047 |
| Parameter | 1047 |
| Command Mode | 1047 |
| Applicability | 1047 |
| Examples | 1047 |
| RSVP-TE Show Commands | 1048 |
| show rsvp label-pool | 1049 |
| Command Syntax | 1049 |
| Parameter | 1049 |
| Command Mode | 1049 |
| Applicability | 1049 |
| Example | 1049 |
| show debugging rsvp | 1050 |
| Command Syntax | 1050 |
| Parameters | 1050 |
| Command Mode | 1050 |
| Applicability | 1050 |
| Example | 1050 |
| show rsvp | 1051 |
| Command Syntax | 1051 |
| Parameters | 1051 |
| Command Mode | 1051 |
| Applicability | 1051 |
| Example | 1051 |
| show rsvp admin-groups | 1053 |
| Command Syntax | 1053 |
| Parameters | 1053 |
| Command Mode | 1053 |
| Applicability | 1053 |
| Example | 1053 |
| show rsvp bypass | 1054 |
| Command Syntax | 1054 |

| | |
|---|------|
| Parameters | 1054 |
| Command Mode | 1054 |
| Applicability | 1054 |
| Example | 1054 |
| show rsvp bypass detail | 1055 |
| Command Syntax | 1055 |
| Parameters | 1055 |
| Command Mode | 1055 |
| Applicability | 1055 |
| Example | 1055 |
| show rsvp bypass lsp-address-list | 1057 |
| Command Syntax | 1057 |
| Parameters | 1057 |
| Command Mode | 1057 |
| Applicability | 1057 |
| Example | 1057 |
| show rsvp bypass protected-lsp-list | 1058 |
| Command Syntax | 1058 |
| Parameters | 1058 |
| Command Mode | 1058 |
| Applicability | 1058 |
| Example | 1058 |
| show rsvp control-adjacency | 1060 |
| Command Syntax | 1060 |
| Parameters | 1060 |
| Command Mode | 1060 |
| Applicability | 1060 |
| Example | 1060 |
| show rsvp data-link | 1062 |
| Command Syntax | 1062 |
| Parameters | 1062 |
| Command Mode | 1062 |
| Applicability | 1062 |
| Example | 1062 |
| show rsvp graceful-restart | 1063 |
| Command Syntax | 1063 |
| Parameters | 1063 |
| Command Mode | 1063 |
| Applicability | 1063 |
| Example | 1063 |
| show rsvp interface | 1064 |
| Command Syntax | 1064 |
| Parameter | 1064 |
| Command Mode | 1064 |
| Applicability | 1064 |
| Example | 1064 |
| show rsvp l2-info | 1066 |

| | |
|---|------|
| Command Syntax | 1066 |
| Parameters | 1066 |
| Command Mode | 1066 |
| Applicability | 1066 |
| Example | 1066 |
| show rsvp local-addresses | 1067 |
| Command Syntax | 1067 |
| Parameters | 1067 |
| Command Mode | 1067 |
| Applicability | 1067 |
| Example | 1067 |
| show rsvp lsp-id-table | 1068 |
| Command Syntax | 1068 |
| Parameters | 1068 |
| Command Mode | 1068 |
| Applicability | 1068 |
| Example | 1068 |
| show rsvp neighbor | 1070 |
| Command Syntax | 1070 |
| Parameters | 1070 |
| Command Mode | 1070 |
| Applicability | 1070 |
| Example | 1070 |
| show rsvp nexthop-cache | 1072 |
| Command Syntax | 1072 |
| Parameters | 1072 |
| Command Mode | 1072 |
| Applicability | 1072 |
| Example | 1072 |
| show rsvp path | 1073 |
| Command Syntax | 1073 |
| Parameter | 1073 |
| Command Mode | 1073 |
| Applicability | 1073 |
| Example | 1073 |
| show rsvp protected-lsp-reop-list | 1074 |
| Command Syntax | 1074 |
| Parameters | 1074 |
| Command Mode | 1074 |
| Applicability | 1074 |
| Example | 1074 |
| show rsvp session | 1075 |
| Command Syntax | 1075 |
| Parameters | 1075 |
| Command Mode | 1075 |
| Applicability | 1075 |
| Example | 1075 |

| | |
|----------------------------------|------|
| show rsvp session count | 1080 |
| Command Syntax | 1080 |
| Parameters | 1080 |
| Command Mode | 1080 |
| Applicability | 1080 |
| Example | 1080 |
| show rsvp session egress | 1081 |
| Command Syntax | 1081 |
| Parameters | 1081 |
| Command Mode | 1081 |
| Applicability | 1081 |
| Examples | 1081 |
| show rsvp session ingress | 1085 |
| Command Syntax | 1085 |
| Parameters | 1085 |
| Command Mode | 1085 |
| Applicability | 1085 |
| Examples | 1085 |
| show rsvp session LSP-NAME | 1089 |
| Command Syntax | 1089 |
| Parameters | 1089 |
| Command Mode | 1089 |
| Applicability | 1089 |
| Example | 1089 |
| show rsvp session transit | 1092 |
| Command Syntax | 1092 |
| Parameters | 1092 |
| Command Mode | 1092 |
| Applicability | 1092 |
| Example | 1092 |
| show rsvp statistics | 1095 |
| Command Syntax | 1095 |
| Parameters | 1095 |
| Command Mode | 1095 |
| Applicability | 1095 |
| Example | 1095 |
| show rsvp summary-refresh | 1096 |
| Command Syntax | 1096 |
| Parameters | 1096 |
| Command Mode | 1096 |
| Applicability | 1096 |
| Example | 1096 |
| show rsvp trunk | 1097 |
| Command Syntax | 1097 |
| Parameters | 1097 |
| Command Mode | 1097 |
| Applicability | 1097 |

| | |
|---|-------------|
| Example | 1097 |
| show rsvp trunk multi-sec-detail | 1099 |
| Command Syntax | 1099 |
| Parameters | 1099 |
| Command Mode | 1099 |
| Applicability | 1099 |
| Example | 1099 |
| show rsvp trunk-id-table | 1100 |
| Command Syntax | 1100 |
| Parameters | 1100 |
| Command Mode | 1100 |
| Applicability | 1100 |
| Example | 1100 |
| show rsvp version | 1102 |
| Command Syntax | 1102 |
| Parameters | 1102 |
| Command Mode | 1102 |
| Applicability | 1102 |
| Example | 1102 |
| Multi-Protocol Label Switching Configuration | 1103 |
| Understanding Label Space | 1106 |
| Overview | 1106 |
| Topology | 1106 |
| Configuration | 1106 |
| PE1 | 1106 |
| P | 1107 |
| PE2 | 1108 |
| Validation | 1109 |
| Understanding MPLS TTL Processing | 1112 |
| Overview | 1112 |
| Topology | 1112 |
| Configuration | 1113 |
| PE1 | 1113 |
| P | 1113 |
| PE2 | 1114 |
| Validation | 1115 |
| MPLS-TE Shared Risk Link Group | 1117 |
| Configure MPLS-TE Shared Risk Link Group | 1117 |
| RTR3 | 1121 |
| Loopback Interface | 1121 |
| MPLS Interfaces | 1121 |
| OSPF | 1122 |
| Global RSVP | 1122 |
| RSVP Interface | 1122 |
| Global SRLG (Path Type: Forced) | 1122 |
| SRLG Interface | 1123 |

| | |
|--|------|
| RTR4 | 1123 |
| Loopback Interface | 1123 |
| MPLS Interface | 1123 |
| OSPF | 1124 |
| Global RSVP | 1124 |
| RSVP Interface | 1124 |
| Global SRLG (Path Type: Forced) | 1124 |
| SRLG Interface | 1124 |
| RTR5 | 1125 |
| Loopback Interface | 1125 |
| MPLS Interface | 1125 |
| OSPF | 1125 |
| Global RSVP | 1126 |
| RSVP Interface | 1126 |
| Global SRLG (Path Type: Forced) | 1126 |
| SRLG Interface | 1126 |
| RTR6 | 1127 |
| Loopback Interface | 1127 |
| MPLS Interface | 1127 |
| OSPF | 1127 |
| Global RSVP | 1128 |
| RSVP Interface | 1128 |
| Global SRLG (Path Type: Forced) | 1128 |
| SRLG Interface | 1128 |
| Validation | 1129 |
| R1 Configuration (Ingress) | 1129 |
| RTR3 Configuration (Transit) | 1136 |
| RTR4 Configuration (Egress) | 1139 |
| RTR5 Configuration | 1143 |
| RTR6 Configuration | 1145 |
| RSVP SRLG with ISIS IGP | 1147 |
| Overview | 1147 |
| Characteristics | 1147 |
| SRLG for Secondary Sessions | 1147 |
| SRLG for Fast Reroute (FRR) | 1147 |
| Benefits | 1148 |
| SRLG with ISIS IGPConfiguration | 1148 |
| Prerequisites | 1148 |
| Topology | 1148 |
| Configuration Steps | 1149 |
| SRLG with ISIS IGP Running Configuration | 1158 |
| Validation | 1165 |
| SRLG with ISIS IGP Commands | 1193 |
| Implementation Examples | 1194 |
| SRLG with One-to-One FRR Protection | 1194 |
| SRLG with Secondary Protection | 1194 |
| SRLG with Secondary-Priority Protection | 1195 |

| | |
|--|------|
| SRLG with Manual Bypass Protection | 1195 |
| SRLG with Auto Bypass Protection | 1196 |
| Troubleshooting | 1196 |
| Glossary | 1196 |
| MPLS DiffServ Configuration | 1198 |
| MPLS Diff-Serv Overview | 1198 |
| Terminology | 1199 |
| Policy Map | 1200 |
| Remarking Configuration | 1201 |
| Configuration | 1201 |
| Topology | 1201 |
| OSPF and LDP Configuration for PE1, P and PE2 | 1201 |
| Configuration of Marking or Remarking | 1204 |
| PE2 | 1207 |
| PE1 | 1208 |
| PE2 | 1211 |
| Policing Configuration | 1219 |
| Configuration | 1219 |
| Topology | 1219 |
| Validation | 1224 |
| MPLS Statistics Configuration | 1230 |
| Topology | 1230 |
| Configure LDP-LSP | 1230 |
| RTR1: Loopback Interface configuration | 1230 |
| RTR1: Global LDP configuration | 1230 |
| RTR1: Enabling LDP and label switching on interface | 1231 |
| RTR1: Global OSPF configuration | 1231 |
| RTR2: Loopback Interface configuration | 1231 |
| RTR2: Global LDP configuration | 1231 |
| RTR2: Enabling LDP and label switching on interface | 1231 |
| RTR2: OSPF Configuration | 1232 |
| RTR3: Loopback Interface configuration | 1232 |
| RTR3: Global LDP configuration | 1232 |
| RTR3: Enabling LDP and label switching on interface | 1232 |
| RTR3: OSPF Configuration | 1233 |
| RTR4: Loopback Interface configuration | 1233 |
| RTR4: Global LDP configuration | 1233 |
| RTR4: Enabling LDP and label switching on interface | 1233 |
| RTR4: Global OSPF Configuration | 1234 |
| MPLS LDP VPLS Configuration | 1234 |
| Virtual Circuit Configuration | 1235 |
| VPLS Configuration | 1236 |
| Configure RSVP-LSP | 1237 |
| RTR1: Global RSVP configuration | 1237 |
| RTR1: Enabling RSVP and label switching on interface | 1237 |
| RTR1: Trunk Configuration | 1237 |
| RTR2: Global RSVP configuration | 1237 |

| | |
|---|------|
| RTR2: Enabling RSVP and label switching on interface | 1237 |
| RTR3: Global RSVP configuration | 1237 |
| RTR3: Enabling RSVP and label switching on interface | 1238 |
| RTR4: Global RSVP configuration | 1238 |
| RTR4: Enabling RSVP and label switching on interface | 1238 |
| RTR4: Trunk Configuration | 1238 |
| Configure Static-LSP | 1238 |
| RTR1: Global Static configuration | 1238 |
| RTR1: Enabling label switching on interface | 1239 |
| RTR2: Global Static configuration | 1239 |
| RTR2: Enabling label switching on interface | 1239 |
| RTR3: Global Static configuration | 1239 |
| RTR3: Enabling label switching on interface | 1239 |
| RTR4: Global Static configuration | 1239 |
| RTR4: Enabling label switching on interface | 1240 |
| Validation | 1240 |
| For Static-LSP | 1240 |
| For RSVP-LSP | 1241 |
| For LDP-LSP | 1242 |
| For LDP-VC | 1243 |
| For LDP-VPLS | 1243 |
| MPLS Service Mapping Configuration | 1245 |
| Overview | 1245 |
| Topology | 1245 |
| Configuration | 1247 |
| PE1: Loopback Interface | 1247 |
| PE1: Global EVPN MPLS Command: | 1247 |
| PE1: LDP Configuration | 1247 |
| PE1: Interface Configuration Network Side | 1248 |
| PE1: OSPF Configuration | 1248 |
| PE1: BGP Configuration | 1249 |
| PE1: SR policy Configuration | 1249 |
| PE1: Tunnel Policy Configuration | 1250 |
| PE1: VPWS Configuration | 1250 |
| PE1: VPLS Configuration | 1251 |
| PE1: MAC VRF and IP VRF Configuration for EVPN and L3VPN services | 1251 |
| PE1: EVPN and MAC-VRF mapping for ELAN and ELINE | 1252 |
| PE1: Access side Configuration | 1252 |
| P1: Loopback Interface | 1254 |
| P1: Interface Configuration Network Side | 1254 |
| P1: OSPF Configuration | 1254 |
| P2: Loopback Interface | 1255 |
| P2: Interface Configuration Network Side | 1255 |
| P2: OSPF Configuration | 1255 |
| PE3: Loopback Interface | 1256 |
| PE3: Global EVPN MPLS Command | 1256 |
| PE3: LDP Configuration | 1257 |

| | |
|--|------|
| PE3: Interface Configuration Network Side | 1257 |
| PE3: OSPF Configuration | 1257 |
| PE3: BGP Configuration | 1258 |
| PE3: VPWS Configuration | 1258 |
| PE3: VPLS Configuration | 1259 |
| PE3: MAC and IP VRF Configuration | 1259 |
| PE3: EVPN and mac vrf mapping Configuration | 1260 |
| PE3: Access side Configuration | 1260 |
| Validation | 1262 |
| PE1 | 1262 |
| VPWS | 1262 |
| VPLS | 1262 |
| EVPN E-LAN | 1263 |
| EVPN E-LINE | 1263 |
| L3VPN and 6VPE | 1264 |
| Tunnel policy Output after applying under all services | 1265 |
| PE2 Validation | 1266 |
| Mapping RSVP Tunnel Name to L2VPN Service | 1269 |
| Topology | 1269 |
| Configuring the Virtual Circuit (VC) | 1269 |
| Configure IP Address and OSPF on Routers | 1270 |
| Configure MPLS, RSVP, and LDP Targeted Peer on Routers | 1271 |
| Configure VC | 1273 |
| Bind Customer Interface to VC | 1274 |
| Validation | 1275 |
| Configuring a MPLS Static Layer-2 VC | 1277 |
| PE1 | 1278 |
| PE2 | 1278 |
| Validation | 1279 |
| Configure Dynamic VPLS | 1282 |
| PE1: LDP VPLS Configuration | 1282 |
| PE2: LDP VPLS Configuration | 1283 |
| LDP VPLS Service Mapping Configuration | 1283 |
| PE1 | 1283 |
| PE1: Access port Configuration | 1284 |
| PE2 | 1284 |
| PE2: Access port Configuration | 1284 |
| Validation | 1284 |
| Configure Static VPLS | 1286 |
| PE1: LDP VPLS Configuration | 1286 |
| PE1: Access port Configuration | 1286 |
| PE2: LDP VPLS Configuration | 1287 |
| PE2: Access port Configuration | 1287 |
| Validation | 1287 |
| MPLS BFD Configuration | 1289 |
| Topology | 1289 |
| Configuring MPLS BFD LSP over LDP | 1289 |

| | |
|---|------|
| PE1 | 1289 |
| P | 1290 |
| PE2 | 1291 |
| Validation | 1292 |
| Configuring MPLS BFD LSP over RSVP | 1293 |
| PE1 | 1293 |
| P | 1294 |
| PE2 | 1294 |
| Validation | 1295 |
| Configuring Static MPLS BFD | 1296 |
| PE1 | 1296 |
| P | 1296 |
| PE2 | 1297 |
| Validations | 1297 |
| Configuring PW VCCV BFD | 1297 |
| Validation | 1300 |
| BGP Peer Groups for Address-Family L2VPN EVPN | 1301 |
| Topology | 1301 |
| Configuration | 1301 |
| R1 | 1301 |
| R2 | 1302 |
| R3 | 1303 |
| Validation | 1304 |
| R1 | 1304 |
| Point-to-Point Connections Over MPLS | 1308 |
| Overview | 1308 |
| Topology | 1308 |
| Configure IP Address and OSPF on Routers | 1308 |
| PE1 | 1309 |
| P1 | 1309 |
| P2 | 1310 |
| PE2 | 1310 |
| PE3 | 1311 |
| Configure MPLS, LDP, and LDP Targeted Peer on Routers | 1311 |
| PE1 | 1311 |
| PE2 | 1312 |
| PE3 | 1312 |
| P1 | 1312 |
| P2 | 1313 |
| Configure VC | 1313 |
| PE1 | 1313 |
| PE2 | 1313 |
| PE3 | 1314 |
| Bind Customer Interface to VC | 1314 |
| PE1 | 1314 |
| PE2 | 1314 |
| PE3 | 1315 |

| | |
|---|-------------|
| Validation | 1315 |
| L2VPN-L3VPN Stitching | 1318 |
| Overview | 1318 |
| Characteristics | 1318 |
| Benefits | 1318 |
| Limitation | 1318 |
| Configuration | 1319 |
| Topology | 1319 |
| PE1 | 1320 |
| P1 | 1321 |
| PE2 | 1321 |
| PE3 | 1325 |
| P2 | 1329 |
| PE4 | 1331 |
| Running Configuration | 1335 |
| Validation | 1347 |
| New CLI | 1373 |
| mpls l2vpn-stitching | 1373 |
| Implementation Examples | 1374 |
| Troubleshooting | 1375 |
| Glossary | 1375 |
| Multi-Protocol Label Switching Command Reference | 1376 |
| MPLS Commands | 1381 |
| admin-groups | 1384 |
| Command Syntax | 1384 |
| Parameters | 1384 |
| Default | 1384 |
| Command Mode | 1384 |
| Applicability | 1384 |
| Example | 1384 |
| advanced parser enable | 1385 |
| Command Syntax | 1385 |
| Parameters | 1385 |
| Default | 1385 |
| Command Mode | 1385 |
| Applicability | 1385 |
| Example | 1385 |
| bandwidth | 1386 |
| Command Syntax | 1386 |
| Parameter | 1386 |
| Default | 1386 |
| Command Mode | 1386 |
| Applicability | 1386 |
| Example | 1386 |
| clear mpls counters ldp | 1387 |
| Command Syntax | 1387 |

| | |
|--|------|
| Parameter | 1387 |
| Command Mode | 1387 |
| Applicability | 1387 |
| Examples | 1387 |
| clear mpls counters rsvp | 1388 |
| Command Syntax | 1388 |
| Parameter | 1388 |
| Command Mode | 1388 |
| Applicability | 1388 |
| Examples | 1388 |
| clear mpls counters static | 1389 |
| Command Syntax | 1389 |
| Parameter | 1389 |
| Command Mode | 1389 |
| Applicability | 1389 |
| Examples | 1389 |
| clear mpls l2-circuit statistics | 1390 |
| Command Syntax | 1390 |
| Parameters | 1390 |
| Command Mode | 1390 |
| Applicability | 1390 |
| Example | 1390 |
| control-word | 1391 |
| Command Syntax | 1391 |
| Parameters | 1391 |
| Default | 1391 |
| Command Mode | 1391 |
| Applicability | 1391 |
| Example | 1391 |
| flow-label | 1392 |
| Command Syntax | 1392 |
| Parameters | 1392 |
| Command Mode | 1392 |
| Applicability | 1392 |
| Example | 1392 |
| group-id | 1394 |
| Command Syntax | 1394 |
| Parameters | 1394 |
| Default | 1394 |
| Command Mode | 1394 |
| Applicability | 1394 |
| Example | 1394 |
| group-name | 1395 |
| Command Syntax | 1395 |
| Parameters | 1395 |
| Default | 1395 |
| Command Mode | 1395 |

| | |
|-----------------------------------|------|
| Applicability | 1395 |
| Example | 1395 |
| label-switching | 1396 |
| Command Syntax | 1396 |
| Parameter | 1396 |
| Default | 1396 |
| Command Mode | 1396 |
| Applicability | 1396 |
| Example | 1396 |
| manual-pseudowire | 1397 |
| Command Syntax | 1397 |
| Parameters | 1397 |
| Default | 1397 |
| Command Mode | 1397 |
| Applicability | 1397 |
| Example | 1397 |
| match vlan | 1398 |
| Command Syntax | 1398 |
| Parameter | 1398 |
| Command Mode | 1398 |
| Applicability | 1398 |
| Example | 1398 |
| mpls ac-group | 1400 |
| Command Syntax | 1400 |
| Parameter | 1400 |
| Default | 1400 |
| Command Mode | 1400 |
| Applicability | 1400 |
| Examples | 1400 |
| mpls bandwidth-class | 1401 |
| Command Syntax | 1401 |
| Parameter | 1401 |
| Default | 1401 |
| Command Mode | 1401 |
| Applicability | 1401 |
| Examples | 1401 |
| mpls-bos parsing nibble-map | 1402 |
| Command Syntax | 1402 |
| Parameters | 1402 |
| Default | 1402 |
| Command Mode | 1402 |
| Applicability | 1402 |
| Example | 1402 |
| mpls ftn-ecmp ldp | 1403 |
| Command Syntax | 1403 |
| Parameter | 1403 |
| Default | 1403 |

| | |
|--------------------------------|------|
| Command Mode | 1403 |
| Applicability | 1403 |
| Examples | 1403 |
| mpls ftn-entry tunnel-id | 1404 |
| Command Syntax | 1404 |
| Parameters | 1404 |
| Default | 1404 |
| Command mode | 1404 |
| Applicability | 1405 |
| Examples | 1405 |
| mpls ftn-entry | 1406 |
| Command Syntax | 1406 |
| Parameters | 1406 |
| Default | 1406 |
| Command Mode | 1406 |
| Applicability | 1406 |
| Example | 1406 |
| mpls ilm-ecmp ldp | 1407 |
| Command Syntax | 1407 |
| Parameters | 1407 |
| Default | 1407 |
| Command Mode | 1407 |
| Applicability | 1407 |
| Example | 1407 |
| mpls ilm-entry pop | 1408 |
| Command Syntax | 1408 |
| Parameters | 1408 |
| Default | 1408 |
| Command Mode | 1408 |
| Applicability | 1408 |
| Examples | 1408 |
| mpls ilm-entry swap | 1409 |
| Command Syntax | 1409 |
| Parameters | 1409 |
| Default | 1409 |
| Command Mode | 1410 |
| Applicability | 1410 |
| Examples | 1410 |
| mpls ingress-ttl | 1411 |
| Command Syntax | 1411 |
| Parameter | 1411 |
| Default | 1411 |
| Command Mode | 1411 |
| Applicability | 1411 |
| Example | 1411 |
| mpls l2-circuit | 1412 |
| Command Syntax | 1412 |

| | |
|------------------------------------|------|
| Parameters | 1412 |
| Default | 1412 |
| Command Mode | 1412 |
| Applicability | 1412 |
| Example | 1412 |
| mpls l2-circuit-fib-entry | 1413 |
| Command Syntax | 1413 |
| Parameters | 1413 |
| Default | 1413 |
| Command Mode | 1413 |
| Applicability | 1413 |
| Example | 1413 |
| mpls l2vpn-nw-statistics | 1414 |
| Command Syntax | 1414 |
| Parameters | 1414 |
| Command Mode | 1414 |
| Default | 1414 |
| Applicability | 1414 |
| Example | 1414 |
| mpls label mode | 1415 |
| Command Syntax | 1415 |
| Parameters | 1415 |
| Default | 1415 |
| Command Mode | 1416 |
| Applicability | 1416 |
| Example | 1416 |
| mpls local-packet-handling | 1417 |
| Command Syntax | 1417 |
| Default | 1417 |
| Parameters | 1417 |
| Command Mode | 1417 |
| Applicability | 1417 |
| Examples | 1417 |
| mpls lsp-encap-dscp-preserve | 1418 |
| Command Syntax | 1418 |
| Default | 1418 |
| Parameters | 1418 |
| Command Mode | 1418 |
| Applicability | 1418 |
| Examples | 1418 |
| mpls lsp-model uniform | 1419 |
| Command Syntax | 1419 |
| Parameter | 1419 |
| Default | 1419 |
| Command Mode | 1419 |
| Applicability | 1419 |
| Examples | 1419 |

| | |
|-----------------------------|------|
| mpls lsp-stitching | 1420 |
| Command Syntax | 1420 |
| Parameters | 1420 |
| Default | 1420 |
| Command Mode | 1420 |
| Applicability | 1420 |
| Command Example | 1420 |
| mpls map-route | 1421 |
| Command Syntax | 1421 |
| Parameters | 1421 |
| Default | 1421 |
| Command Mode | 1421 |
| Applicability | 1421 |
| Examples | 1421 |
| mpls min-label-value | 1422 |
| Command Syntax | 1422 |
| Parameters | 1422 |
| Default | 1422 |
| Command Mode | 1423 |
| Applicability | 1423 |
| Examples | 1423 |
| mpls propagate-ttl | 1424 |
| Command Syntax | 1424 |
| Parameters | 1424 |
| Default | 1424 |
| Command Mode | 1424 |
| Applicability | 1424 |
| Examples | 1424 |
| mpls traffic-eng srlg | 1425 |
| Command Syntax: | 1425 |
| Parameters | 1425 |
| Command Mode | 1425 |
| Example | 1425 |
| pce entity | 1426 |
| Syntax | 1426 |
| Parameters | 1426 |
| Command Mode | 1426 |
| Applicability | 1426 |
| Example | 1426 |
| pce lsp-delegate | 1427 |
| Syntax | 1427 |
| Parameters | 1427 |
| Default | 1427 |
| Command Mode | 1427 |
| Applicability | 1427 |
| Example | 1427 |
| pce state-report | 1428 |

| | |
|--|------|
| Syntax | 1428 |
| Parameters | 1428 |
| Default | 1428 |
| Command Mode | 1428 |
| Applicability | 1428 |
| Example | 1428 |
| ping mpls | 1429 |
| Command Syntax | 1429 |
| Parameters | 1429 |
| Defaults | 1431 |
| Command Mode | 1431 |
| Applicability | 1431 |
| Example | 1431 |
| rewrite ingress | 1433 |
| Command Syntax | 1433 |
| Parameters | 1433 |
| Command Mode | 1433 |
| Applicability | 1433 |
| Examples | 1433 |
| secondary-priority srlg-disjoint | 1434 |
| Command Syntax | 1434 |
| Parameters | 1434 |
| Command Mode | 1434 |
| Example | 1434 |
| secondary srlg-disjoint | 1435 |
| Command Syntax | 1435 |
| Parameters | 1435 |
| Command Mode | 1435 |
| Example | 1435 |
| service-tpid | 1436 |
| Command Syntax | 1436 |
| Parameters | 1436 |
| Default | 1436 |
| Command Mode | 1436 |
| Applicability | 1436 |
| Example | 1436 |
| show admin-groups | 1437 |
| Command Syntax | 1437 |
| Parameters | 1437 |
| Command Mode | 1437 |
| Applicability | 1437 |
| Examples | 1437 |
| show hsl mpls l3vpn-ftn | 1438 |
| Command Syntax | 1438 |
| Parameters | 1438 |
| Command Mode | 1438 |
| Applicability | 1438 |

| | |
|--|------|
| Examples | 1438 |
| show mpls bandwidth-class | 1439 |
| Command Syntax | 1439 |
| Parameters | 1439 |
| Command Mode | 1439 |
| Applicability | 1439 |
| Examples | 1439 |
| show mpls counters ldp | 1440 |
| Command Syntax | 1440 |
| Parameter | 1440 |
| Command Mode | 1440 |
| Applicability | 1440 |
| Examples | 1440 |
| show mpls | 1442 |
| Command Syntax | 1442 |
| Parameters | 1442 |
| Command Mode | 1442 |
| Applicability | 1442 |
| Examples | 1442 |
| show mpls counters static | 1444 |
| Command Syntax | 1444 |
| Parameter | 1444 |
| Command Mode | 1444 |
| Applicability | 1444 |
| Examples | 1444 |
| show mpls counters rsvp | 1446 |
| Command Syntax | 1446 |
| Parameter | 1446 |
| Command Mode | 1446 |
| Applicability | 1446 |
| Examples | 1446 |
| show mpls cross-connect-table | 1448 |
| Command Syntax | 1448 |
| Parameters | 1448 |
| Command Mode | 1448 |
| Applicability | 1448 |
| Example | 1448 |
| show mpls cross-connect-table backup | 1450 |
| Command Syntax | 1450 |
| Parameter | 1450 |
| Command Mode | 1450 |
| Applicability | 1450 |
| Example | 1450 |
| show mpls forwarding-table detail | 1451 |
| Command Syntax | 1451 |
| Parameters | 1451 |
| Command Mode | 1451 |

| | |
|---------------------------------------|------|
| Applicability | 1451 |
| Example | 1451 |
| show mpls forwarding-table | 1453 |
| Command Syntax | 1453 |
| Parameters | 1453 |
| Command Mode | 1453 |
| Applicability | 1453 |
| Example | 1453 |
| show mpls ftn-table | 1456 |
| Command Syntax | 1456 |
| Parameters | 1456 |
| Command Mode | 1456 |
| Applicability | 1456 |
| Example | 1456 |
| show mpls ilm-table | 1458 |
| Command Syntax | 1458 |
| Parameters | 1458 |
| Command Mode | 1458 |
| Applicability | 1458 |
| Example | 1458 |
| show mpls in-segment-table | 1461 |
| Command Syntax | 1461 |
| Parameters | 1461 |
| Command Mode | 1461 |
| Applicability | 1461 |
| Example | 1461 |
| show mpls l2-circuit | 1463 |
| Command Syntax | 1463 |
| Parameters | 1463 |
| Command Mode | 1463 |
| Applicability | 1463 |
| Example | 1463 |
| show mpls l2-circuit statistics | 1465 |
| Command Syntax | 1465 |
| Parameters | 1465 |
| Command Mode | 1465 |
| Applicability | 1465 |
| Example | 1465 |
| show mpls mapped-routes | 1467 |
| Command Syntax | 1467 |
| Parameters | 1467 |
| Command Mode | 1467 |
| Applicability | 1467 |
| Example | 1467 |
| show mpls out-segment-table | 1468 |
| Command Syntax | 1468 |
| Parameters | 1468 |

| | |
|--|------|
| Command Mode | 1468 |
| Applicability | 1468 |
| Example | 1468 |
| show mpls owner-rib-table | 1470 |
| Command Syntax | 1470 |
| Parameters | 1470 |
| Example | 1470 |
| show mpls qos-resource | 1471 |
| Command Syntax | 1471 |
| Parameters | 1471 |
| Command Mode | 1471 |
| Applicability | 1471 |
| Example | 1471 |
| show mpls vc-table | 1473 |
| Command Syntax | 1473 |
| Parameters | 1473 |
| Command Mode | 1473 |
| Applicability | 1473 |
| Examples | 1473 |
| show mpls vrf | 1474 |
| Command Syntax | 1474 |
| Parameters | 1474 |
| Command Mode | 1474 |
| Applicability | 1474 |
| Examples | 1474 |
| show mpls vrf-forwarding-table vrf | 1475 |
| Command Syntax | 1475 |
| Parameters | 1475 |
| Command Mode | 1475 |
| Applicability | 1475 |
| Examples | 1475 |
| show running-config interface mpls | 1478 |
| Command Syntax | 1478 |
| Parameters | 1478 |
| Command Mode | 1478 |
| Applicability | 1478 |
| Example | 1478 |
| show running-config mpls | 1479 |
| Command Syntax | 1479 |
| Parameters | 1479 |
| Command Mode | 1479 |
| Applicability | 1479 |
| Example | 1479 |
| show running-config vc | 1480 |
| Command Syntax | 1480 |
| Parameters | 1480 |
| Command Mode | 1480 |

| | |
|------------------------------------|------|
| Applicability | 1480 |
| Examples | 1480 |
| show running-config vpls | 1481 |
| Command Syntax | 1481 |
| Parameters | 1481 |
| Command Mode | 1481 |
| Applicability | 1481 |
| Examples | 1481 |
| show vccv statistics | 1482 |
| Command Syntax | 1482 |
| Parameters | 1482 |
| Command Mode | 1482 |
| Applicability | 1482 |
| Examples | 1482 |
| srlg-disjoint | 1483 |
| Command Syntax | 1483 |
| Parameters | 1483 |
| Default | 1483 |
| Command Mode | 1483 |
| Applicability | 1484 |
| Example | 1484 |
| trace mpls | 1485 |
| Command Syntax | 1485 |
| Parameters | 1485 |
| Default | 1485 |
| Command Mode | 1486 |
| Applicability | 1486 |
| Examples | 1486 |
| suppress-oper-log mpls l2vpn | 1487 |
| Command Syntax | 1487 |
| Parameters | 1487 |
| Default | 1487 |
| Command Mode | 1487 |
| Applicability | 1487 |
| Example | 1487 |
| trace mpls | 1488 |
| Command Syntax | 1488 |
| Parameters | 1488 |
| Default | 1489 |
| Command Mode | 1489 |
| Applicability | 1489 |
| Examples | 1489 |
| tunnel-name | 1490 |
| Command Syntax | 1490 |
| Parameters | 1490 |
| Default | 1490 |
| Command Mode | 1490 |

| | |
|--|------|
| Applicability | 1490 |
| Example | 1490 |
| tunnel-id | 1491 |
| Command Syntax | 1491 |
| Parameters | 1491 |
| Default | 1491 |
| Command Mode | 1491 |
| Applicability | 1491 |
| Example | 1491 |
| tunnel-select-policy | 1492 |
| Command Syntax | 1492 |
| Parameters | 1492 |
| Default | 1492 |
| Command Mode | 1492 |
| Applicability | 1492 |
| Example | 1492 |
| vccv cc-type | 1493 |
| Command Syntax | 1493 |
| Parameters | 1493 |
| Default | 1493 |
| Command Mode | 1493 |
| Applicability | 1493 |
| Example | 1493 |
| vccv cv-type | 1494 |
| Command Syntax | 1494 |
| Parameters | 1494 |
| Default | 1494 |
| Command Mode | 1494 |
| Applicability | 1494 |
| Example | 1494 |
| Virtual Private LAN Service Commands | 1495 |
| ac-admin-status | 1496 |
| Command Syntax | 1496 |
| Parameter | 1496 |
| Default | 1496 |
| Command Mode | 1496 |
| Applicability | 1496 |
| Example | 1496 |
| ac-description | 1497 |
| Command Syntax | 1497 |
| Parameter | 1497 |
| Default | 1497 |
| Command Mode | 1497 |
| Applicability | 1497 |
| Example | 1497 |
| clear mpls vpls | 1498 |
| Command Syntax | 1498 |

| | |
|--|------|
| Parameters | 1498 |
| Command Mode | 1498 |
| Applicability | 1498 |
| Example | 1498 |
| control-word | 1499 |
| Command Syntax | 1499 |
| Parameters | 1499 |
| Default | 1499 |
| Command Mode | 1499 |
| Applicability | 1499 |
| Example | 1499 |
| exit-signaling | 1500 |
| Command Syntax | 1500 |
| Parameters | 1500 |
| Default | 1500 |
| Command Mode | 1500 |
| Applicability | 1500 |
| Examples | 1500 |
| exit-if-vpls | 1501 |
| Command Syntax | 1501 |
| Parameter | 1501 |
| Default | 1501 |
| Command Mode | 1501 |
| Applicability | 1501 |
| Example | 1501 |
| learning disable (VPLS Mode) | 1502 |
| Command Syntax | 1502 |
| Parameter | 1502 |
| Default | 1502 |
| Command Mode | 1502 |
| Applicability | 1502 |
| Example | 1502 |
| learning disable (Interface VPLS Mode) | 1503 |
| Command Syntax | 1503 |
| Parameter | 1503 |
| Default | 1503 |
| Command Mode | 1503 |
| Applicability | 1503 |
| Example | 1503 |
| learning enable | 1504 |
| Command Syntax | 1504 |
| Parameter | 1504 |
| Default | 1504 |
| Command Mode | 1504 |
| Applicability | 1504 |
| Example | 1504 |
| mpls vpls | 1505 |

| | |
|------------------------------------|------|
| Command Syntax | 1505 |
| Parameters | 1505 |
| Default | 1505 |
| Command Mode | 1505 |
| Applicability | 1505 |
| Example | 1505 |
| mpls-vpls service-template | 1506 |
| Command Syntax | 1506 |
| Parameters | 1506 |
| Default | 1506 |
| Command Mode | 1506 |
| Applicability | 1506 |
| Example | 1506 |
| no learning | 1507 |
| Command Syntax | 1507 |
| Parameter | 1507 |
| Default | 1507 |
| Command Mode | 1507 |
| Applicability | 1507 |
| Example | 1507 |
| show mpls vpls | 1508 |
| Command Syntax | 1508 |
| Parameters | 1508 |
| Command Mode | 1508 |
| Applicability | 1508 |
| Examples | 1508 |
| show mpls vpls mac-address | 1516 |
| Command Syntax | 1516 |
| Parameters | 1516 |
| Command Mode | 1516 |
| Applicability | 1516 |
| Examples | 1516 |
| show mpls vpls split-horizon | 1518 |
| Command Syntax | 1518 |
| Parameters | 1518 |
| Default | 1518 |
| Command Mode | 1518 |
| Applicability | 1518 |
| Example | 1518 |
| show mpls vpls statistics | 1520 |
| Command Syntax | 1520 |
| Parameters | 1520 |
| Command Mode | 1520 |
| Applicability | 1520 |
| Example | 1520 |
| signaling ldp | 1522 |
| Command Syntax | 1522 |

| | |
|----------------------------|------|
| Parameters | 1522 |
| Default | 1522 |
| Command Mode | 1522 |
| Applicability | 1522 |
| Examples | 1522 |
| signaling bgp | 1523 |
| Command Syntax | 1523 |
| Parameters | 1523 |
| Default | 1523 |
| Command Mode | 1523 |
| Applicability | 1523 |
| Examples | 1523 |
| static-mac | 1524 |
| Command Syntax | 1524 |
| Parameter | 1524 |
| Default | 1524 |
| Command Mode | 1524 |
| Applicability | 1524 |
| Examples | 1524 |
| tunnel-select-policy | 1525 |
| Command Syntax | 1525 |
| Parameters | 1525 |
| Command Mode | 1525 |
| Applicability | 1525 |
| Examples | 1525 |
| ve-id | 1526 |
| Command Syntax | 1526 |
| Parameters | 1526 |
| Default | 1526 |
| Command Mode | 1526 |
| Applicability | 1526 |
| Examples | 1526 |
| vpls-ac-group | 1527 |
| Command Syntax | 1527 |
| Parameter | 1527 |
| Default | 1527 |
| Command Mode | 1527 |
| Applicability | 1527 |
| Examples | 1527 |
| vpls-description | 1528 |
| Command Syntax | 1528 |
| Parameter | 1528 |
| Default | 1528 |
| Command Mode | 1528 |
| Applicability | 1528 |
| Example | 1528 |
| vpls fib-entry | 1529 |

| | |
|------------------------|------|
| Command Syntax | 1529 |
| Parameters | 1529 |
| Default | 1529 |
| Command Mode | 1529 |
| Applicability | 1529 |
| Examples | 1529 |
| vpls-mtu | 1531 |
| Command Syntax | 1531 |
| Parameter | 1531 |
| Default | 1531 |
| Command Mode | 1531 |
| Applicability | 1531 |
| Example | 1531 |
| vpls-peer | 1532 |
| Command Syntax | 1532 |
| Parameters | 1532 |
| Default | 1533 |
| Command Mode | 1533 |
| Applicability | 1533 |
| Examples | 1533 |
| vpls-peer manual | 1534 |
| Command Syntax | 1534 |
| Parameters | 1534 |
| Default | 1534 |
| Command Mode | 1534 |
| Applicability | 1534 |
| Examples | 1534 |
| vpls-type | 1535 |
| Command Syntax | 1535 |
| Parameter | 1535 |
| Default | 1535 |
| Command Mode | 1535 |
| Applicability | 1535 |
| Examples | 1535 |
| vpls-vc | 1536 |
| Command Syntax | 1536 |
| Parameter | 1536 |
| Default | 1536 |
| Command Mode | 1536 |
| Applicability | 1536 |
| Example | 1536 |
| vpls-route-map | 1537 |
| Command Syntax | 1537 |
| Parameter | 1537 |
| Default | 1537 |
| Command Mode | 1537 |
| Applicability | 1537 |

| | |
|--------------------------------|------|
| Example | 1537 |
| EVPN MPLS Commands | 1538 |
| access-if-evpn | 1540 |
| Command Syntax | 1540 |
| Parameters | 1540 |
| Command Mode | 1540 |
| Applicability | 1540 |
| Examples | 1540 |
| arp-cache disable | 1541 |
| Command Syntax | 1541 |
| Parameters | 1541 |
| Default | 1541 |
| Command Mode | 1541 |
| Applicability | 1541 |
| Examples | 1541 |
| arp-nd cos | 1542 |
| Command Syntax | 1542 |
| Parameters | 1542 |
| Default | 1542 |
| Command Mode | 1542 |
| Applicability | 1542 |
| Examples | 1542 |
| arp-nd flood-suppress | 1543 |
| Command Syntax | 1543 |
| Parameters | 1543 |
| Default | 1543 |
| Command Mode | 1543 |
| Applicability | 1543 |
| Examples | 1543 |
| arp-nd refresh timer | 1544 |
| Command Syntax | 1544 |
| Parameters | 1544 |
| Command Mode | 1544 |
| Applicability | 1544 |
| Examples | 1544 |
| clear evpn mpls counters | 1545 |
| Command Syntax | 1545 |
| Parameters | 1545 |
| Command Mode | 1545 |
| Applicability | 1545 |
| Examples | 1545 |
| clear mac address-table | 1546 |
| Command Syntax | 1546 |
| Parameters | 1546 |
| Command Mode | 1546 |
| Applicability | 1546 |
| Example | 1546 |

| | |
|--|------|
| dynamic-learning disable | 1547 |
| Command Syntax | 1547 |
| Parameters | 1547 |
| Default | 1547 |
| Command Mode | 1547 |
| Applicability | 1547 |
| Example | 1547 |
| evi-name | 1548 |
| Command Syntax | 1548 |
| Parameters | 1548 |
| Default | 1548 |
| Command Mode | 1548 |
| Applicability | 1548 |
| Example | 1548 |
| evpn esi hold-time | 1549 |
| Command Syntax | 1549 |
| Parameters | 1549 |
| Default | 1549 |
| Command Mode | 1549 |
| Applicability | 1549 |
| Example | 1549 |
| evpn etree | 1550 |
| Command Syntax | 1550 |
| Parameters | 1550 |
| Default | 1550 |
| Command Mode | 1550 |
| Applicability | 1550 |
| Example | 1550 |
| evpn irb | 1551 |
| Command Syntax | 1551 |
| Parameters | 1551 |
| Command Mode | 1551 |
| Applicability | 1551 |
| Example | 1551 |
| evpn irb-forwarding anycast-gateway-mac | 1552 |
| Command Syntax | 1552 |
| Parameters | 1552 |
| Command Mode | 1552 |
| Applicability | 1552 |
| Example | 1552 |
| evpn irb-if-forwarding anycast-gateway-mac | 1553 |
| Command Syntax | 1553 |
| Parameters | 1553 |
| Command Mode | 1553 |
| Applicability | 1553 |
| Example | 1553 |
| evpn mpls enable | 1554 |

| | |
|------------------------------------|------|
| Command Syntax | 1554 |
| Parameters | 1554 |
| Command Mode | 1554 |
| Applicability | 1554 |
| Examples | 1554 |
| evpn mpls id | 1555 |
| Command Syntax | 1555 |
| Parameters | 1555 |
| Command Mode | 1555 |
| Applicability | 1555 |
| Example | 1555 |
| evpn mpls irb | 1557 |
| Command Syntax | 1557 |
| Parameters | 1557 |
| Command Mode | 1557 |
| Applicability | 1557 |
| Example | 1557 |
| l3vni | 1558 |
| Command Syntax | 1558 |
| Parameters | 1558 |
| Command Mode | 1558 |
| Applicability | 1558 |
| Example | 1558 |
| mac | 1559 |
| Command Syntax | 1559 |
| Parameters | 1559 |
| Command Mode | 1559 |
| Applicability | 1559 |
| Example | 1559 |
| evpn mpls mac-ageing-time | 1560 |
| Command Syntax | 1560 |
| Parameters | 1560 |
| Default | 1560 |
| Command Mode | 1560 |
| Applicability | 1560 |
| Example | 1560 |
| evpn mpls multihoming enable | 1561 |
| Command Syntax | 1561 |
| Parameters | 1561 |
| Default | 1561 |
| Command Mode | 1561 |
| Applicability | 1561 |
| Example | 1561 |
| evpn mpls vtep-ip-global | 1562 |
| Command Syntax | 1562 |
| Parameters | 1562 |
| Command Mode | 1562 |

| | |
|---|------|
| Applicability | 1562 |
| Examples | 1562 |
| evpn multi-homed | 1563 |
| Command Syntax | 1563 |
| Parameters | 1563 |
| Default | 1563 |
| Command Mode | 1563 |
| Applicability | 1563 |
| Example | 1563 |
| hardware-profile filter evpn-mpls-cw | 1565 |
| Command Syntax | 1565 |
| Parameters | 1565 |
| Default | 1565 |
| Command Mode | 1565 |
| Applicability | 1565 |
| Example | 1565 |
| hardware-profile filter evpn-mpls-mh | 1566 |
| Command Syntax | 1566 |
| Parameters | 1566 |
| Default | 1566 |
| Command Mode | 1566 |
| Applicability | 1566 |
| Example | 1566 |
| hardware-profile statistics evpn-mpls | 1567 |
| Command Syntax | 1567 |
| Parameters | 1567 |
| Default | 1567 |
| Command Mode | 1567 |
| Applicability | 1567 |
| Example | 1567 |
| host-reachability-protocol | 1568 |
| Command Syntax | 1568 |
| Parameters | 1568 |
| Default | 1568 |
| Command Mode | 1568 |
| Applicability | 1568 |
| Examples | 1568 |
| evpn-route-map | 1569 |
| Command Syntax | 1569 |
| Parameters | 1569 |
| Default | 1569 |
| Command Mode | 1569 |
| Applicability | 1569 |
| Examples | 1569 |
| ip address | 1570 |
| Command Syntax | 1570 |
| Default | 1570 |

| | |
|-------------------------------|------|
| Parameters | 1570 |
| Command Mode | 1570 |
| Applicability | 1570 |
| Example | 1570 |
| ipv6 address | 1571 |
| Command Syntax | 1571 |
| Default | 1571 |
| Parameters | 1571 |
| Command Mode | 1571 |
| Applicability | 1571 |
| Example | 1571 |
| link-loss forwarding | 1572 |
| Command Syntax | 1572 |
| Parameters | 1572 |
| Default | 1572 |
| Command Mode | 1572 |
| Applicability | 1572 |
| Example | 1572 |
| mac-holdtime | 1573 |
| Command Syntax | 1573 |
| Parameters | 1573 |
| Default | 1573 |
| Command Mode | 1573 |
| Applicability | 1573 |
| Examples | 1573 |
| mac-vrf | 1574 |
| Command Syntax | 1574 |
| Parameters | 1574 |
| Command Mode | 1574 |
| Applicability | 1574 |
| Examples | 1574 |
| map vpn-id | 1575 |
| Command Syntax | 1575 |
| Parameters | 1575 |
| Command Mode | 1575 |
| Applicability | 1575 |
| Examples | 1575 |
| nd-cache disable | 1576 |
| Command Syntax | 1576 |
| Parameters | 1576 |
| Default | 1576 |
| Command Mode | 1576 |
| Applicability | 1576 |
| Examples | 1576 |
| rewrite egr-pcp-disable | 1577 |
| Command Syntax | 1577 |
| Parameters | 1577 |

| | |
|---------------------------------|------|
| Default | 1577 |
| Command Mode | 1577 |
| Applicability | 1577 |
| Example | 1577 |
| service-carving | 1578 |
| Command Syntax | 1578 |
| Parameters | 1578 |
| Default | 1578 |
| Command Mode | 1578 |
| Applicability | 1578 |
| Example | 1578 |
| service-carving ac-driven | 1580 |
| Command Syntax | 1580 |
| Parameters | 1580 |
| Default | 1580 |
| Command Mode | 1580 |
| Applicability | 1580 |
| Example | 1580 |
| service-carving weight | 1581 |
| Command Syntax | 1581 |
| Parameters | 1581 |
| Default | 1581 |
| Command Mode | 1581 |
| Applicability | 1581 |
| Example | 1581 |
| show bgp l2vpn evpn | 1582 |
| Command Syntax | 1582 |
| Parameters | 1582 |
| Command Mode | 1582 |
| Applicability | 1582 |
| Examples | 1582 |
| show evpn load-balance | 1584 |
| Command Syntax | 1584 |
| Parameters | 1584 |
| Command Mode | 1584 |
| Applicability | 1584 |
| Example | 1584 |
| show evpn mpls | 1586 |
| Command Syntax | 1586 |
| Parameters | 1586 |
| Command Mode | 1586 |
| Applicability | 1586 |
| Example | 1586 |
| show evpn mpls arp-cache | 1587 |
| Command syntax | 1587 |
| Parameters | 1587 |
| Command Mode | 1587 |

| | |
|--|------|
| Applicability | 1587 |
| Examples | 1587 |
| show evpn mpls counters | 1588 |
| Command syntax | 1588 |
| Parameters | 1588 |
| Command Mode | 1588 |
| Applicability | 1588 |
| Examples | 1588 |
| show evpn mpls label alias | 1589 |
| Command syntax | 1589 |
| Parameters | 1589 |
| Command Mode | 1589 |
| Applicability | 1589 |
| Examples | 1589 |
| show evpn mpls label esi | 1590 |
| Command syntax | 1590 |
| Parameters | 1590 |
| Command Mode | 1590 |
| Applicability | 1590 |
| Examples | 1590 |
| show evpn mpls mac-table | 1591 |
| Command Syntax | 1591 |
| Parameters | 1591 |
| Command Mode | 1591 |
| Applicability | 1591 |
| Example | 1591 |
| show evpn mpls nd-cache | 1593 |
| Command Syntax | 1593 |
| Parameters | 1593 |
| Command Mode | 1593 |
| Applicability | 1593 |
| Examples | 1593 |
| show evpn mpls route-count | 1594 |
| Command Syntax | 1594 |
| Parameters | 1594 |
| Command Mode | 1594 |
| Applicability | 1594 |
| Examples | 1594 |
| show evpn mpls static host state | 1596 |
| Command Syntax | 1596 |
| Parameters | 1596 |
| Command Mode | 1596 |
| Applicability | 1596 |
| Examples | 1596 |
| show evpn mpls status | 1597 |
| Command Syntax | 1597 |
| Parameters | 1597 |

| | |
|--|------|
| Command Mode | 1597 |
| Applicability | 1597 |
| Example | 1597 |
| show evpn mpls tunnel | 1598 |
| Command Syntax | 1598 |
| Parameters | 1598 |
| Command Mode | 1598 |
| Applicability | 1598 |
| Example | 1598 |
| show evpn mpls tunnel label | 1599 |
| Command Syntax | 1599 |
| Parameters | 1599 |
| Command Mode | 1599 |
| Applicability | 1599 |
| Example | 1599 |
| show evpn mpls xconnect | 1600 |
| Command Syntax | 1600 |
| Parameters | 1600 |
| Command Mode | 1600 |
| Applicability | 1600 |
| Example | 1600 |
| show evpn mpls xconnect tunnel | 1601 |
| Command Syntax | 1601 |
| Parameters | 1601 |
| Command Mode | 1601 |
| Applicability | 1601 |
| Example | 1601 |
| show evpn mpls xconnect tunnel label | 1602 |
| Command Syntax | 1602 |
| Parameters | 1602 |
| Command Mode | 1602 |
| Applicability | 1602 |
| Example | 1602 |
| show running-config evpn mpls | 1603 |
| Command Syntax | 1603 |
| Parameters | 1603 |
| Command Mode | 1603 |
| Applicability | 1603 |
| Examples | 1603 |
| tunnel-select-policy | 1605 |
| Command Syntax | 1605 |
| Parameters | 1605 |
| Command Mode | 1605 |
| Applicability | 1605 |
| Examples | 1605 |
| MPLS Service Mapping Commands | 1606 |
| color | 1607 |

| | |
|--|------|
| Command Syntax | 1607 |
| Parameters | 1607 |
| Command Mode | 1607 |
| Applicability | 1607 |
| Example | 1607 |
| mode | 1608 |
| Command Syntax | 1608 |
| Parameters | 1608 |
| Default | 1608 |
| Command Mode | 1608 |
| Applicability | 1608 |
| Examples | 1608 |
| mpls 6pe tunnel-select-policy | 1609 |
| Command Syntax | 1609 |
| Parameters | 1609 |
| Command Mode | 1609 |
| Applicability | 1609 |
| Examples | 1609 |
| show running-config tunnel-policy | 1610 |
| Command Syntax | 1610 |
| Parameters | 1610 |
| Command Mode | 1610 |
| Applicability | 1610 |
| Examples | 1610 |
| show tunnel-policy | 1611 |
| Command Syntax | 1611 |
| Parameters | 1611 |
| Command Mode | 1611 |
| Applicability | 1611 |
| Examples | 1611 |
| tunnel-policy | 1614 |
| Command Syntax | 1614 |
| Parameters | 1614 |
| Command Mode | 1614 |
| Applicability | 1614 |
| Example | 1614 |
| tunnel-select-policy (For L3VPN service) | 1615 |
| Command Syntax | 1615 |
| Parameters | 1615 |
| Command Mode | 1615 |
| Applicability | 1615 |
| Examples | 1615 |
| MPLS LSP BFD Commands | 1616 |
| lsp ping-interval | 1617 |
| Command Syntax | 1617 |
| Parameters | 1617 |
| Command Mode | 1617 |

| | |
|---|------|
| Applicability | 1617 |
| Example | 1617 |
| min-tx | 1618 |
| Command Syntax | 1618 |
| Parameters | 1618 |
| Defaults | 1618 |
| Command Mode | 1618 |
| Applicability | 1618 |
| Example | 1618 |
| mpls bfd | 1619 |
| Command Syntax | 1619 |
| Parameters | 1619 |
| Command Mode | 1619 |
| Applicability | 1619 |
| Example | 1619 |
| shutdown | 1620 |
| Command Syntax | 1620 |
| Parameters | 1620 |
| Command Mode | 1620 |
| Applicability | 1620 |
| Example | 1620 |
| L2VPN over L2 Subinterface Commands | 1621 |
| access-if-vpws | 1622 |
| Command Syntax | 1622 |
| Parameters | 1622 |
| Command Mode | 1622 |
| Applicability | 1622 |
| Examples | 1622 |
| mpls-l2-circuit | 1623 |
| Command Syntax | 1623 |
| Parameters | 1623 |
| Command Mode | 1623 |
| Applicability | 1623 |
| Examples | 1623 |
| vc-mode | 1624 |
| Command Syntax | 1624 |
| Parameters | 1624 |
| Command Mode | 1624 |
| Applicability | 1624 |
| Examples | 1624 |
| access-if-vpls | 1625 |
| Command Syntax | 1625 |
| Parameters | 1625 |
| Command Mode | 1625 |
| Applicability | 1625 |
| Examples | 1625 |
| mpls-vpls | 1626 |

| | |
|---|-------------|
| Command Syntax | 1626 |
| Parameters | 1626 |
| Command Mode | 1626 |
| Applicability | 1626 |
| Examples | 1626 |
| learning disable | 1627 |
| Command Syntax | 1627 |
| Parameters | 1627 |
| Command Mode | 1627 |
| Applicability | 1627 |
| Examples | 1627 |
| split-horizon (subinterface) | 1628 |
| Command Syntax | 1628 |
| Parameters | 1628 |
| Default | 1628 |
| Command Mode | 1628 |
| Applicability | 1628 |
| Examples | 1628 |
| static-mac | 1629 |
| Command Syntax | 1629 |
| Parameters | 1629 |
| Command Mode | 1629 |
| Applicability | 1629 |
| Examples | 1629 |
| Ethernet Virtual Private Network Configuration | 1630 |
| EVPN MPLS Configuration | 1633 |
| Overview | 1633 |
| VPN Terminology | 1633 |
| EVPN MPLS Single Homing | 1634 |
| Topology | 1634 |
| Validation | 1642 |
| PE1: E-LAN or EVPN-VPLS | 1642 |
| PE1: E-LINE or EVPN-VPWS | 1643 |
| PE2: ELAN or EVPN-VPLS | 1644 |
| PE2: ELINE or EVPN-VPWS | 1645 |
| EVPN MPLS Single Homing Over BGP-LU | 1646 |
| Topology | 1646 |
| Validation | 1657 |
| PE1: EVPN-VPLS or ELAN | 1657 |
| PE1: EVPN-VPWS or ELINE | 1659 |
| PE2: EVPN-VPLS or ELAN | 1660 |
| PE2: ELINE or EVPN-VPWS | 1661 |
| EVPN MPLS Multihoming | 1662 |
| Topology | 1662 |
| Validation | 1674 |
| EVPN MPLS Multihoming over BGP-LU | 1680 |

| | |
|--|------|
| Topology | 1680 |
| Validation | 1696 |
| PE1: ELAN or EVPN-VPLS | 1696 |
| PE1: ELINE or EVPN-VPWS | 1698 |
| PE2: ELAN or EVPN-VPLS | 1699 |
| PE2: ELINE or EVPN-VPWS | 1700 |
| PE3: ELAN or EVPN-VPLS: | 1701 |
| PE3: ELINE or EVPN-VPWS | 1702 |
| EVPN-MPLS L2CP Tunneling | 1703 |
| Topology | 1703 |
| EVPN-MPLS MAC Statistics | 1704 |
| Topology | 1705 |
| EVPN-MPLS Control-Word Support | 1705 |
| Topology | 1705 |
| EVPN-MPLS RSVP Multipath | 1708 |
| Topology | 1708 |
| EVPN-ELINE CFM Single Homing | 1711 |
| Topology | 1711 |
| Validation | 1719 |
| LDP ECMP Configuration | 1722 |
| Topology | 1722 |
| Configurations | 1722 |
| Validation | 1737 |
| EVPN Active-Standby | 1752 |
| Overview | 1752 |
| Single-Active | 1752 |
| Port-Active | 1752 |
| Feature Characteristics | 1752 |
| Benefits | 1754 |
| Prerequisites | 1754 |
| Configuration | 1755 |
| EVPN MPLS Active-Standby MultiHoming Configuration | 1755 |
| EVPN MPLS Active-Standby MH Topology | 1755 |
| Customer Edge (CE) Routers Configurations | 1755 |
| PE1 Router Configurations | 1757 |
| PE2 Router Configurations | 1763 |
| PE3 Router Configurations | 1770 |
| PE4 Router Configurations | 1776 |
| Core Routers Configurations | 1782 |
| EVPN MPLS Active-Standby MH Validation | 1786 |
| EVPN SR Active-Standby Multi-Homing Configuration | 1798 |
| EVPN SR Active-Standby MH Topology | 1798 |
| Customer Edge (CE) Routers Configurations | 1799 |
| Provider Edge Routers Configurations | 1801 |
| Core Routers Configurations | 1830 |
| EVPN SR Active-Standby MH Validation | 1836 |
| Implementation Examples | 1852 |

| | |
|---|------|
| Revised CLI Commands | 1852 |
| evpn multi-homed | 1852 |
| CLI Commands | 1852 |
| service-carving ac-driven | 1853 |
| service-carving | 1854 |
| service-carving weight | 1855 |
| Glossary | 1856 |
| EVPN MPLS E-Tree Scenario 1 | 1859 |
| Overview | 1859 |
| Feature Characteristics | 1859 |
| Benefits | 1860 |
| Prerequisites | 1860 |
| Configuration | 1863 |
| Topology | 1863 |
| Validation | 1865 |
| Static MAC-IP Advertisement | 1867 |
| Network Topology Running Configurations | 1869 |
| E-Tree Active-Standby Configuration | 1881 |
| Implementation Examples | 1885 |
| E-Tree Scenario 1 Commands | 1886 |
| Revised Commands | 1886 |
| Troubleshooting | 1886 |
| Glossary | 1887 |
| EVPN MPLS E-Tree Scenario 2 | 1888 |
| Overview | 1888 |
| Inter-PE E-tree: New Route Advertisements | 1888 |
| Feature Characteristics | 1888 |
| Benefits | 1890 |
| E-Tree Scenario 2 Prerequisites | 1890 |
| Root PEs and Leaf PEs | 1890 |
| MPLS Core Fabric Formation (P1 or P2 Routers) | 1892 |
| Access Switch VLAN Trunk Configuration | 1893 |
| E-Tree Scenario 2 Configuration | 1893 |
| Topology | 1893 |
| Enable E-Tree Scenario-2 Functionality | 1894 |
| VPN ID and Leaf Role Mapping | 1895 |
| E-Tree Scenario 2 Validation | 1895 |
| E-Tree Scenario 2 Network Topology Running Configurations | 1899 |
| E-Tree Scenario 2 Implementation Examples | 1918 |
| E-Tree Scenario 2 Troubleshooting | 1918 |
| E-Tree Scenario 2 Revised Commands | 1919 |
| evpn etree | 1920 |
| map vpn-id | 1921 |
| E-Tree Scenario 2 Glossary | 1921 |
| Flow-Aware Transport for EVPN over MPLS | 1924 |
| Overview | 1924 |
| Feature Characteristics | 1924 |

| | |
|--|------|
| Benefits | 1924 |
| FAT Prerequisites | 1924 |
| PE Nodes | 1925 |
| FAT Configuration | 1928 |
| Topology | 1928 |
| PE Nodes Configuration | 1930 |
| FAT Validation | 1931 |
| FAT Implementation Example | 1934 |
| FAT Revised Command | 1935 |
| evpn mpls id | 1936 |
| FAT Glossary | 1937 |
| EVPN MAC Limit Enforcement at EVI and AC for EVPN-MPLS | 1939 |
| Overview | 1939 |
| Feature Characteristics | 1939 |
| Benefits | 1939 |
| Prerequisites | 1939 |
| Limitations | 1940 |
| Topology | 1940 |
| Configuration | 1940 |
| EVPN prerequisite configurations: | 1940 |
| EVPN-MAC limit profile configuration: | 1941 |
| Configuration snapshot | 1941 |
| Validation | 1941 |
| CLI Commands | 1942 |
| mac-limit-profile | 1942 |
| learning-limit | 1943 |
| action | 1944 |
| high-watermark | 1944 |
| low-watermark | 1945 |
| errdisable-timeout | 1946 |
| learning limit | 1946 |
| Anycast Gateway Routing for Multiple Subnets in EVPN-IRB | 1949 |
| Overview | 1949 |
| Feature Characteristics | 1949 |
| Benefits | 1949 |
| Topology | 1950 |
| Configuration | 1950 |
| PE1 Configuration | 1950 |
| P Configuration | 1954 |
| PE2 Configuration | 1955 |
| PE2 Configuration | 1959 |
| Validation | 1963 |
| Abbreviations | 1981 |
| CFM over EVPN-MPLS for ELINE MultiHoming | 1982 |
| Overview | 1982 |
| Feature Characteristics | 1982 |
| Continuity Check Message | 1982 |

| | |
|--|-------------|
| Ping and Trace | 1982 |
| Benefits | 1983 |
| Prerequisites | 1983 |
| Configuration | 1984 |
| Validation | 1998 |
| Implementation Examples | 2000 |
| Troubleshooting | 2000 |
| Abbreviations | 2001 |
| Glossary | 2001 |
| EVPN EPL Link-Loss Forwarding | 2003 |
| Overview | 2003 |
| Feature Characteristics | 2003 |
| Benefits | 2003 |
| Configuration of LLF for an EVPN Ethernet Private Line (EPL) | 2003 |
| Topology | 2003 |
| Configuration XE22.2 | 2004 |
| CLI Commands | 2005 |
| llf-enable | 2006 |
| Command Syntax | 2006 |
| Parameters | 2006 |
| Default | 2006 |
| Command Mode | 2006 |
| Applicability | 2006 |
| Example | 2006 |
| Glossary | 2006 |
| TWAMP over EVPN Configuration | 2007 |
| Topology | 2007 |
| Configure TWAMP over EVPN for LDP | 2008 |
| EVPN-Eline configurations | 2008 |
| Validation | 2013 |
| Layer 2 Virtual Private Network Configuration | 2016 |
| Virtual Private Wire Service Configuration | 2019 |
| Overview | 2019 |
| Configuration | 2019 |
| Configure IP Address and OSPF on Routers | 2019 |
| Configure MPLS, LDP, and LDP Targeted Peer on Routers | 2021 |
| Configure VC | 2022 |
| Bind Customer Interface to VC | 2023 |
| Layer 2 Tagged Traffic | 2024 |
| Validation | 2025 |
| Configure a Static Layer-2 VC | 2027 |
| Validation | 2028 |
| Sub-Interface with Multiple Match | 2028 |
| PE-1 | 2028 |
| PE-2 | 2029 |
| Validation | 2029 |

| | |
|--|------|
| VPWS-CFM | 2030 |
| Prerequisite | 2030 |
| PE1 | 2031 |
| PE2 | 2031 |
| Validation | 2032 |
| Virtual Private LAN Service Configuration | 2034 |
| VPLS Raw Mode | 2034 |
| VPLS Mesh Peer Configuration | 2034 |
| PE1 | 2034 |
| P | 2035 |
| PE2 | 2037 |
| PE3 | 2038 |
| Validation | 2039 |
| PE1 | 2039 |
| P | 2040 |
| PE3 | 2040 |
| VPLS Tagged Mode | 2042 |
| Configuration | 2043 |
| PE1 | 2043 |
| P | 2045 |
| PE2 | 2046 |
| PE3 | 2048 |
| Validation | 2049 |
| Validation for the Number of Configured VPLS Instances | 2053 |
| VPLS Split Horizon Group | 2053 |
| Configuration With Service Template | 2054 |
| PE1 | 2054 |
| P1 | 2058 |
| PE2 | 2059 |
| P2 | 2063 |
| Validation | 2064 |
| VPLS - Per-Peer VPLS-type Signaling | 2068 |
| Overview | 2068 |
| Feature Characteristics | 2068 |
| Configuration | 2069 |
| PE1 Configuration | 2069 |
| P Configuration | 2070 |
| PE2 Configuration | 2071 |
| PE3 Configuration | 2071 |
| Validation | 2072 |
| VPLS - Tunnel-ID | 2077 |
| Overview | 2077 |
| Feature Characteristics | 2077 |
| Configuration | 2078 |
| Validation | 2081 |
| VPLS - Tunnel-Name | 2087 |
| Overview | 2087 |

| | |
|--|------|
| Feature Characteristics | 2087 |
| Configuration | 2088 |
| Validation | 2091 |
| L2VPN FAT SUPPORT | 2100 |
| Topology | 2100 |
| Configuration for VPLS with FAT | 2100 |
| Validation | 2108 |
| PE1 | 2108 |
| PE2 | 2109 |
| P1 | 2110 |
| P4 | 2110 |
| Configuration for BGP VPLS with FAT | 2111 |
| Validation | 2119 |
| PE1 | 2119 |
| PE2 | 2120 |
| P1 | 2121 |
| Configuration for VPWS with FAT | 2122 |
| Validation | 2129 |
| PE1 | 2129 |
| PE2 | 2130 |
| P1 | 2130 |
| P4 | 2131 |
| L2VPN Sub Interface Configuration | 2132 |
| Topology | 2132 |
| VPLS Configuration | 2132 |
| Validation | 2135 |
| VPWS Configuration | 2136 |
| Validation | 2139 |
| Static VPLS Configuration | 2141 |
| Overview | 2141 |
| Configure Static VPLS | 2141 |
| Validation | 2146 |
| Remove Configurations | 2148 |
| BGP-VPLS Configuration | 2149 |
| Topology | 2149 |
| BGP-VPLS Configuration | 2150 |
| Validation | 2156 |
| BGP Four Bytes ASN with VPLS Configuration | 2162 |
| Overview | 2162 |
| Prerequisites | 2163 |
| Topology | 2163 |
| Configuration | 2164 |
| Validation | 2167 |
| Static VPLS Service Mapping Configuration | 2174 |
| Overview | 2174 |
| Configuration | 2174 |
| Topology | 2174 |

| | |
|--|------|
| PE1: Loopback Interface | 2174 |
| PE1: Interface Configuration | 2175 |
| PE1: Static VPLS Configuration | 2175 |
| PE1: OSPF Configuration | 2175 |
| PE1: FIB Entry Configuration | 2176 |
| P: Loopback Interface | 2176 |
| P: Interface Configuration | 2176 |
| P: FIB Entry Configuration | 2177 |
| PE2: Loopback Interface | 2177 |
| PE2: Interface Configuration | 2178 |
| OSPF Configuration | 2178 |
| PE2: Static VPLS Configuration | 2178 |
| PE2: FIB Entry Configuration | 2179 |
| Configuration of SubInterface | 2179 |
| PE1: Access port Configuration | 2179 |
| PE1: POP | 2179 |
| PE1: XLATE | 2180 |
| PE1: PUSH | 2180 |
| PE2: Access port Configuration | 2180 |
| PE2: POP | 2181 |
| PE2: XLATE | 2181 |
| PE2: PUSH | 2181 |
| Validation | 2182 |
| LDP-VPLS Service Mapping Configuration | 2186 |
| Overview | 2186 |
| Topology | 2186 |
| Configuration | 2186 |
| PE1: Loopback Interface | 2186 |
| PE1: Global LDP | 2186 |
| PE1: Interface Configuration | 2187 |
| PE1: OSPF Configuration | 2187 |
| PE1: LDP VPLS Configuration | 2187 |
| P: Loopback Interface | 2188 |
| P: Global LDP | 2188 |
| P: Interface Configuration | 2188 |
| P: OSPF Configuration | 2189 |
| PE2: Loopback Interface | 2189 |
| PE2: Global LDP | 2189 |
| PE2: Interface Configuration | 2189 |
| PE2: OSPF Configuration | 2189 |
| PE2: LDP VPLS Configuration | 2190 |
| LDP VPLS Subinterface Matching | 2191 |
| PE1: POP | 2191 |
| PE1: XLATE | 2191 |
| PE1: PUSH | 2191 |
| PE1:POP-2tag | 2192 |
| PE2: POP | 2192 |

| | |
|---|------|
| PE2: XLATE | 2192 |
| PE2: PUSH | 2192 |
| PE2:POP-2tag | 2193 |
| Validation | 2193 |
| BGP-VPLS Service Mapping Configuration | 2199 |
| Overview | 2199 |
| Topology | 2199 |
| Configuration | 2199 |
| BGP VPLS SUBINTERFACE ATTACHMENT | 2204 |
| Validation | 2206 |
| Hierarchical VPLS | 2213 |
| Overview | 2213 |
| H-VPLS Redundancy Characteristics | 2213 |
| Benefits | 2213 |
| Limitations | 2213 |
| Prerequisites | 2213 |
| ISIS Configuration | 2214 |
| OSPF Configuration | 2214 |
| Configuration for H-VPLS with Redundancy | 2214 |
| Topology | 2214 |
| Configure H-VPLS on PE1 Router | 2215 |
| Configure H-VPLS on PE2 (Primary Hub) | 2216 |
| Configure H-VPLS on PE3 (Secondary Hub) | 2216 |
| Configure H-VPLS on Spoke Router | 2217 |
| Running Configuration | 2218 |
| Validation | 2219 |
| Configuration for H-VPLS without Redundancy | 2221 |
| Topology | 2221 |
| Configure H-VPLS on PE1 Router | 2221 |
| Configure H-VPLS on Hub Router | 2222 |
| Configure H-VPLS on Spoke Router | 2223 |
| Running Configuration | 2223 |
| Validation | 2224 |
| Commands for H-VPLS Configuration | 2226 |
| vpls-vc | 2226 |
| signaling | 2227 |
| MAC Withdrawal - VPLS/H-VPLS | 2229 |
| Overview | 2229 |
| Characteristics of MAC Withdrawal - VPLS/H-VPLS | 2229 |
| Benefits | 2229 |
| Prerequisites | 2230 |
| Configuration | 2231 |
| Topology | 2231 |
| PE1 | 2231 |
| PE2 (Primary Hub) | 2232 |
| PE3 (Secondary Hub) | 2233 |
| Spoke with VPLS | 2233 |

| | |
|--|------|
| Spoke with VPWS | 2234 |
| Running Configuration | 2235 |
| PE1 | 2235 |
| PE2 | 2235 |
| PE3 | 2236 |
| Spoke with VPLS | 2236 |
| Spoke with VPWS | 2236 |
| Validation | 2237 |
| CLI Commands | 2240 |
| mac-withdrawal flush-on-spoke-vc-standby-activation | 2241 |
| mac-withdrawal flush-propagate-spoke-to-mesh | 2242 |
| mac-withdrawal flush-on-vpws-vc-switchover | 2243 |
| Glossary | 2243 |
| MAC Move Protection - VPLS/H-VPLS | 2244 |
| Overview | 2244 |
| Characteristics of MAC Move Protection - VPLS/H-VPLS | 2244 |
| Benefits | 2244 |
| Prerequisites | 2244 |
| Configuration | 2245 |
| Topology | 2245 |
| PE1 | 2246 |
| Hub | 2246 |
| Spoke | 2247 |
| Running Configuration | 2248 |
| PE1 | 2248 |
| Hub | 2248 |
| Spoke | 2249 |
| Validation | 2249 |
| CLI Commands | 2252 |
| vpls mac-move enable detect | 2253 |
| mac-move | 2254 |
| show mpls vpls mac-move name | 2255 |
| Glossary | 2255 |
| MAC Limit for VPLS and H-VPLS | 2256 |
| Overview | 2256 |
| Characteristics of MAC Move Limit - VPLS/H-VPLS | 2256 |
| Benefits | 2256 |
| Prerequisites | 2256 |
| Configuration | 2257 |
| Topology | 2257 |
| PE1 | 2257 |
| Hub | 2258 |
| Spoke | 2259 |
| Running Configuration | 2260 |
| PE1 | 2260 |
| Hub | 2260 |
| Spoke | 2261 |

| | |
|---|------|
| Validation | 2261 |
| CLI Commands | 2263 |
| vpls mac-limit-profile | 2264 |
| Glossary | 2265 |
| MLAG Active-Standby for VPLS | 2266 |
| Overview | 2266 |
| Feature Characteristics | 2266 |
| Benefit | 2267 |
| Prerequisites | 2267 |
| Configuration | 2267 |
| Topology | 2267 |
| Configuring MLAG on PE1 (Active) and PE2 (Standby) | 2267 |
| Configuring VPLS Session on PE1 (Active MLAG Peer), PE2 (Standby MLAG Peer), and PE3 | 2268 |
| Running configurations | 2269 |
| PE1 | 2269 |
| PE2 | 2269 |
| PE3 | 2270 |
| Validation | 2270 |
| Verify MLAG Domain summary on PE1 and PE2 to verify the Current MLAG status and MLAG Synchronization | 2270 |
| Verify VPLS sessions on PE1, PE2, and PE3 | 2271 |
| Verify VPLS MAC addresses learned on PE1, PE2, and PE3 | 2272 |
| MLAG Active peer went down (Shutdown Active MLAG peer interface on PE1 or on switch) | 2272 |
| Verify MLAG domain summary on PE1 and PE2 to verify the Current MLAG status and MLAG Synchronization | 2273 |
| Verify VPLS sessions on PE1, PE2, and PE3 | 2273 |
| Verify VPLS MAC addresses learned on PE1, PE2, and PE3 | 2274 |
| MLAG Active-Standby for VPLS commands | 2275 |
| switchover type | 2276 |
| mac flush send on mlag switchover | 2277 |
| Glossary | 2277 |
| BGP Auto-Discovery (AD) for LDP VPLS | 2279 |
| Overview | 2279 |
| Benefits | 2279 |
| Limitations | 2279 |
| Prerequisites | 2280 |
| Configuration | 2281 |
| Topology | 2281 |
| PE5 | 2281 |
| P1 | 2283 |
| PE3 | 2284 |
| PE4 | 2285 |
| PE2 | 2287 |
| Running Configuration | 2288 |
| PE5 | 2288 |

| | |
|---|------|
| P1 | 2289 |
| PE3 | 2290 |
| PE4 | 2292 |
| PE2 | 2293 |
| Validation | 2294 |
| PE4 | 2297 |
| PE2 | 2302 |
| CLI Commands | 2304 |
| bgp-auto-discovery | 2305 |
| l2vpn-id | 2306 |
| rd (route distinguisher) | 2307 |
| route-target | 2308 |
| Ethernet Data Plane Loopback (EDPL) | 2310 |
| Overview | 2310 |
| Feature Characteristics | 2310 |
| Benefits | 2310 |
| Prerequisites | 2310 |
| Limitation | 2311 |
| Configuration | 2311 |
| Topology | 2311 |
| Enable EDPL for IPv4 | 2312 |
| Enable EDPL for IPv6 | 2313 |
| Enable EDPL for L2 | 2313 |
| Validation | 2314 |
| Implementation Examples | 2315 |
| EDPL CLI Commands | 2317 |
| address-family | 2318 |
| direction | 2319 |
| destination-mac | 2320 |
| destination-udp-port | 2321 |
| destination-ip-address | 2322 |
| duration | 2323 |
| edpl profile | 2324 |
| edpl start | 2325 |
| edpl stop | 2326 |
| edpl test | 2327 |
| ethertype | 2328 |
| ip-swap | 2329 |
| source-port | 2330 |
| source-ip-address | 2331 |
| source-udp-port | 2332 |
| udp-tcp-port-swap | 2333 |
| vlan | 2334 |
| show ethernet edpl | 2335 |
| show ethernet edpl detail | 2336 |
| Troubleshooting | 2336 |
| Glossary | 2337 |

| | |
|--|-------------|
| Layer 3 Virtual Private Network Configuration | 2339 |
| MPLS Layer-3 VPN Configurations | 2341 |
| Overview | 2341 |
| Terminology | 2341 |
| VPN Routing Process | 2342 |
| Configure MPLS Layer-3 VPN | 2343 |
| Enable IGP | 2345 |
| Enable IGP Switching: PE1 | 2345 |
| Enable IGP Switching: P | 2345 |
| Enable IGP Switching: PE2 | 2345 |
| Enable Label Switching Protocol | 2345 |
| LDP: PE1 | 2346 |
| LDP: P | 2346 |
| LDP: PE2 | 2346 |
| RSVP-TE: PE1 | 2346 |
| RSVP-TE: P | 2347 |
| RSVP-TE: PE2 | 2347 |
| Configure PEs as BGP Neighbors | 2347 |
| PE1 | 2348 |
| PE2 | 2348 |
| PE1 | 2349 |
| PE2 | 2349 |
| PE1 | 2349 |
| PE2 | 2350 |
| Configure VRF—RD and Route Targets | 2350 |
| Configure Route Distinguishers | 2350 |
| PE1 | 2350 |
| PE2 | 2351 |
| Configure Route Targets | 2351 |
| PE1 | 2351 |
| PE2 | 2352 |
| Configure CE Neighbor for the VPN (Using BGP/ OSPF) | 2352 |
| Using BGP | 2352 |
| PE1 | 2352 |
| PE2 | 2353 |
| CE1: BGP | 2353 |
| CE2: BGP | 2354 |
| PE1 | 2354 |
| PE2 | 2355 |
| Verify the MPLS-VPN Configuration | 2355 |
| Inter-AS VPN Configuration Overview | 2359 |
| Types of Inter-AS VPN | 2359 |
| Inter-AS VPN Option-A Overview | 2359 |
| Topology | 2360 |
| Inter-AS VPN Option-A Configuration | 2360 |
| Validation | 2368 |
| Inter-AS VPN Option-B Overview | 2370 |

| | |
|---|------|
| Topology | 2371 |
| Inter-AS VPN Option-B Configuration | 2371 |
| Validation | 2380 |
| Inter-AS VPN Option-C Configuration | 2382 |
| Inter-AS Option-C with RR Configuration | 2383 |
| Validation | 2395 |
| Configuration Without RR | 2403 |
| Validation | 2412 |
| ECMP Support for L3VPN | 2418 |
| Overview | 2418 |
| Prerequisites | 2419 |
| Configuration | 2420 |
| PE1 | 2420 |
| RR1 | 2421 |
| RR2 | 2422 |
| PE2 | 2423 |
| PE3 | 2424 |
| PE1 | 2425 |
| RR1 | 2426 |
| RR2 | 2426 |
| PE2 | 2427 |
| PE3 | 2428 |
| Validation | 2429 |
| Scenarios | 2430 |
| ECMP for L3VPN Commands | 2431 |
| Example 1 | 2432 |
| Example 2 | 2432 |
| ECMP Support for L3EVPN | 2433 |
| Overview | 2433 |
| Prerequisites | 2434 |
| Configuration | 2435 |
| PE1 | 2436 |
| P1/RR1 | 2437 |
| P2 | 2438 |
| PE3 | 2438 |
| PE4 | 2439 |
| Configuration Snapshot | 2440 |
| P1/RR1 | 2453 |
| P2/RR2 | 2459 |
| PE3 | 2463 |
| PE4 | 2477 |
| Validation | 2489 |
| ECMP for L3EVPN Commands | 2496 |
| L3VPN over EVPN MPLS Configuration | 2498 |
| Topology | 2498 |
| Configuration | 2498 |
| Validation | 2514 |

| | |
|--|------|
| L3VPN GR Configuration | 2522 |
| Topology | 2522 |
| Configuration | 2523 |
| Validation | 2527 |
| Restart BGP Gracefully | 2527 |
| After Restarting the BGP Manually | 2528 |
| TWAMP over L3VPN Configuration | 2531 |
| Topology | 2531 |
| Configure TWAMP over L3VPN for LDP | 2532 |
| Validation | 2537 |
| Bridge Virtual Interface (BVI) Over L3VPN | 2540 |
| Overview | 2540 |
| Characteristics of BVI Over L3VPN | 2540 |
| Benefits | 2540 |
| Limitations | 2540 |
| Prerequisites | 2541 |
| Configuration | 2542 |
| Topology | 2542 |
| Configure CE1 Router | 2542 |
| Configure PE1 Router | 2543 |
| Configure PE2 Router | 2544 |
| Configure CE2 Router | 2544 |
| Running Configuration | 2545 |
| CE1 | 2545 |
| PE1 | 2546 |
| P1 | 2548 |
| PE2 | 2550 |
| CE2 | 2552 |
| Validation | 2553 |
| Glossary | 2554 |
| DSCP Preserve | 2555 |
| DSCP Preserve Characteristics | 2555 |
| DSCP Preserve Global | 2556 |
| Prerequisites | 2556 |
| Configuration for DSCP Preserve Global | 2556 |
| Topology | 2556 |
| Configure DSCP Preserve and L3VPN Services on PE1 Router | 2556 |
| Running Configuration on PE1 Router is as follows: | 2558 |
| Running Configuration on P Router is as follows: | 2560 |
| Running Configuration on PE2 Router is as follows: | 2561 |
| Validation | 2563 |
| Commands | 2568 |
| mpls lsp-encap-dscp-preserve | 2568 |
| DSCP Preserve Per-VRF | 2570 |
| DSCP Preserve Per-VRF Characteristics | 2570 |
| Prerequisites | 2570 |
| Configuration for DSCP Preserve Per-VRF | 2571 |

| | |
|---|-------------|
| Topology | 2571 |
| Configure DSCP Preserve and L3VPN Services on PE1 Router | 2571 |
| Running Configuration on PE1 Router is as follows: | 2573 |
| Running Configuration on P Router is as follows: | 2575 |
| Running Configuration on PE2 Router is as follows: | 2576 |
| Validation | 2578 |
| DSCP Commands | 2586 |
| mpls lsp-encap-dscp-preserve | 2587 |
| mpls 6pe lsp-encap-dscp-preserve | 2588 |
| Glossary | 2588 |
| IPv6 Provider Edge Routers (6PE) /MPLS VPN Configuration | 2590 |
| 6PE Configuration | 2591 |
| Benefits of 6PE | 2591 |
| IPv6 on Provider Edge Routers | 2591 |
| Configuring 6PE | 2592 |
| Topology | 2592 |
| Configuration | 2592 |
| Validation | 2597 |
| CE1 | 2597 |
| PE1 | 2597 |
| P1 | 2600 |
| PE2 | 2601 |
| CE2 | 2604 |
| 6VPE Configuration | 2606 |
| Topology | 2606 |
| Configuration | 2606 |
| CE1 | 2606 |
| CE2 | 2607 |
| PE1 | 2607 |
| P1 | 2609 |
| PE2 | 2610 |
| Validation | 2611 |
| CE1 | 2611 |
| PE1 | 2612 |
| P1 | 2616 |
| PE2 | 2616 |
| CE2 | 2620 |
| Bridge Virtual Interface (BVI) Over 6vPE | 2621 |
| Overview | 2621 |
| Prerequisites | 2621 |
| Configuration | 2622 |
| Topology | 2622 |
| Configure CE1 Router | 2623 |
| Configure PE1 Router | 2623 |
| Configure PE2 Router | 2624 |
| Configure CE2 Router | 2624 |

| | |
|--|-------------|
| Running Configuration | 2625 |
| CE1 | 2625 |
| PE1 | 2625 |
| PE2 | 2626 |
| CE2 | 2626 |
| CE2 | 2628 |
| Validation | 2628 |
| Glossary | 2630 |
| Appendix A: Service Template Configuration | 2631 |
| Overview | 2631 |
| Purpose and Function | 2631 |
| Typical Use Cases | 2631 |
| Common Syntax Patterns | 2632 |
| VPLS – Untagged Interface Configuration | 2632 |
| VPLS – VLAN Tagged Interface Configuration | 2632 |
| VPLS – VLAN Tagged Interface with Rewrite | 2632 |
| VPWS – Untagged Interface Configuration | 2633 |
| VPWS – VLAN Tagged Interface Configuration | 2633 |
| VPWS – VLAN Tagged Interface with Rewrite | 2633 |
| Documentation Reference | 2634 |
| Migration from Service Template to Sub-interface Configuration | 2634 |
| Guidance | 2634 |
| Key Capabilities | 2634 |
| Example | 2635 |
| VPWS Sub-interface-based Configuration (Recommended): | 2635 |
| VPWS Service Template-based Configuration: | 2635 |
| VPLS Sub-interface-based Configuration (Recommended): | 2635 |
| VPLS Service Template-based Configuration: | 2635 |

PREFACE

About this Guide

This guide describes how to configure Multi-Protocol Label Switching in OcNOS.


Audience

This guide is intended for network administrators and other engineering professionals who configure OcNOS.

Conventions

The [Table 1](#) table shows the conventions used in this guide.

Table 1. Conventions

| Convention | Description |
|---|--|
| Italics | Emphasized terms; titles of books |
|  Note: | Special instructions, suggestions, or warnings |
| <code>monospaced type</code> | Code elements such as commands, parameters, files, and directories |

IP Infusion Product Release Version

Each integer in release numbers indicates Major, Minor, and Maintenance release versions. Build numbers that follow the release numbers are for internal tracking and verification of the software build process and are visible to customers as part of the software version number.



Product Name: IP Infusion Product Family

Major Version: New customer-facing functionality that represents a significant change to the code base; including a significant marketing change or direction in the product.

Minor Version: Enhancements or extensions to existing features, changes to address external needs, or internal improvements to satisfy new sales regions or marketing initiatives.

Maintenance Version: A collection of product bugs or issues usually scheduled every 30 or 60 days, based on the number of issues.

Related Documentation

For information about installing OcNOS, see the *Installation Guide* for your platform.

Feature Availability

Each OcNOS SKU contains a set of supported features. For a list of available features based on the SKU that you purchased, refer to the [Feature Matrix](#).

Migration Guide

Check the *Migration Guide* for necessary configuration changes before migrating from one version of OcNOS to another.

IP Maestro Support

Monitor devices running OcNOS Release 6.3.4-70 and above using IP Maestro software.

Technical Support

IP Infusion maintains an online technical support site that provides a variety of technical support programs for licensed OcNOS customers at the [Technical Assistance Center](#).

Customers and partners enjoy full access to the support website. The site allows customers and partners to open technical support calls, update open calls with new information, and review the status of open or closed calls. The password-protected site includes technical documentation, Release Notes, and descriptions of service offerings.

Technical Sales

Contact the IP Infusion sales representative for more information about the OcNOS solution.

Technical Documentation

For core commands and configuration procedures, visit: [Product Documentation](#).

For training videos, visit: [OcNOS Free Training Videos](#).

For a list of supported platforms and SKUs of OcNOS features, refer to the [OcNOS Feature Matrix](#).

Documentation Disclaimer

The global documentation site is evolving to provide an enhanced website user experience for select topics included in this release. Some guides are now available outside the existing documentation library and can be accessed directly from custom documentation landing pages. These guides offer robust in-built search functionality.

For the latest documentation, visit the product-specific documentation landing page and select the relevant guide.

Comments

If you have comments, or need to report a problem with the content, contact techpubs@ipinfusion.com.

Command Line Interface

This chapter introduces the OcNOS Command Line Interface (CLI) and how to use its features.

Overview

You use the CLI to configure, monitor, and maintain OcNOS devices. The CLI is text-based and each command is usually associated with a specific task.

You can give the commands described in this manual locally from the console of a device running OcNOS or remotely from a terminal emulator such as putty or xterm. You can also use the commands in scripts to automate configuration tasks.

Chapter Organization

The chapters in command references are organized as described in [Command Description Format \(page 110\)](#).

The chapters in configuration guides are organized into these major sections:

- An overview that explains a configuration in words
 - Topology with a diagram that shows the devices and connections used in the configuration
 - Configuration steps in a table for each device where the left-hand side shows the commands you enter and the right-hand side explains the actions that the commands perform
 - Validation which shows commands and their output that verify the configuration
-

Command Line Interface Help

You access the CLI help by entering a full or partial command string and a question mark “?”. The CLI displays the command keywords or parameters along with a short description. For example, at the CLI command prompt, type:

```
> show ?
```

The CLI displays this keyword list with short descriptions for each keyword:

```
show ?
  application-priority  Application Priority
  arp                  Internet Protocol (IP)
  bfd                  Bidirectional Forwarding Detection (BFD)
  bgp                  Border Gateway Protocol (BGP)
  bi-lsp               Bi-directional lsp status and configuration
  bridge               Bridge group commands
  ce-vlan              COS Preservation for Customer Edge VLAN
  class-map            Class map entry
  cli                  Show CLI tree of current mode
  clns                 Connectionless-Mode Network Service (CLNS)
  control-adjacency    Control Adjacency status and configuration
  control-channel      Control Channel status and configuration
  cspf                 CSPF Information
  customer             Display Customer spanning-tree
  cvlan                Display CVLAN information
  debugging            Debugging functions
  etherchannel         LACP etherchannel
```

```

    ethernet
    ...
    Layer-2

```

If you type the ? in the middle of a keyword, the CLI displays help for that keyword only.

```

> show de?
debugging  Debugging functions

```

If you type the ? in the middle of a keyword, but the incomplete keyword matches several other keywords, OcNOS displays help for all matching keywords.

```

> show i? (CLI does not display the question mark).
interface  Interface status and configuration
ip         IP information
isis      ISIS information

```

Command Completion

The CLI can complete the spelling of a command or a parameter. Begin typing the command or parameter and then press the tab key. For example, at the CLI command prompt type **sh**:

```

> sh

```

Press the tab key. The CLI displays:

```

> show

```

If the spelling of a command or parameter is ambiguous, the CLI displays the choices that match the abbreviation. Type **show i** and press the tab key. The CLI displays:

```

> show i
interface  ip          ipv6      isis
> show i

```

The CLI displays the **interface** and **ip** keywords. Type **n** to select **interface** and press the tab key. The CLI displays:

```

> show in
> show interface

```

Type **?** and the CLI displays the list of parameters for the **show interface** command.

```

> show interface
IFNAME  Interface name
|       Output modifiers
>       Output redirection
<cr>

```

The CLI displays the only parameter associated with this command, the **IFNAME** parameter.

Command Abbreviations

The CLI accepts abbreviations that uniquely identify a keyword in commands. For example:

```

> sh int xe0

```

is an abbreviation for:

```

> show interface xe0

```

Command Line Errors

Any unknown spelling causes the CLI to display the error **Unrecognized command** in response to the ?. The CLI displays the command again as last entered.

```
> show dd?
% Unrecognized command
> show dd
```

When you press the Enter key after typing an invalid command, the CLI displays:

```
(config)#router ospf here
                        ^
% Invalid input detected at '^' marker.
```

where the ^ points to the first character in error in the command.

If a command is incomplete, the CLI displays the following message:

```
> show
% Incomplete command.
```

Some commands are too long for the display line and can wrap mid-parameter or mid-keyword, as shown below. This does *not* cause an error and the command performs as expected:

```
area 10.10.0.18 virtual-link 10.10.0.19 authentication-key 57393
```

Command Negation

Many commands have a **no** form that resets a feature to its default value or disables the feature. For example:

- The **ip address** command assigns an IPv4 address to an interface
- The **no ip address** command removes an IPv4 address from an interface

Syntax Conventions

[Table 2](#) describes the conventions used to represent command syntax in this reference.

Table 2. Syntax conventions

| Convention | Description | Example |
|------------------------|--|---------------------------|
| monospaced font | Command strings entered on a command line | show ip ospf |
| lowercase | Keywords that you enter exactly as shown in the command syntax. | show ip ospf |
| UPPERCASE | See Variable Placeholders (page 109) | IFNAME |
| () | Optional parameters, from which you must select one. Vertical bars delimit the selections. Do not enter the parentheses or vertical bars as part of the command. | (A.B.C.D <0-4294967295>) |
| () | Optional parameters, from which you select one or none. Vertical bars delimit the | (A.B.C.D <0-4294967295>) |

Table 2. Syntax conventions (continued)

| Convention | Description | Example |
|------------|--|--|
| | selections. Do not enter the parentheses or vertical bars as part of the command. | |
| () | Optional parameter which you can specify or omit. Do not enter the parentheses or vertical bar as part of the command. | (IFNAME) |
| { } | Optional parameters, from which you must select one or more. Vertical bars delimit the selections. Do not enter the braces or vertical bars as part of the command. | {intra-area <1-255> inter-area <1-255> external <1-255>} |
| [] | Optional parameters, from which you select zero or more. Vertical bars delimit the selections. Do not enter the brackets or vertical bars as part of the command. | [<1-65535> AA:NN internet local-AS no-advertise no-export] |
| ? | Nonrepeatable parameter. The parameter that follows a question mark can only appear once in a command string. Do not enter the question mark as part of the command. | ?route-map WORD |
| . | Repeatable parameter. The parameter that follows a period can be repeated more than once. Do not enter the period as part of the command. | set as-path prepend .<1-65535> |

Variable Placeholders

[Table 3](#) shows the tokens used in command syntax use to represent variables for which you supply a value.

Table 3. Variable placeholders

| Token | Description |
|------------|--|
| WORD | A contiguous text string (excluding spaces) |
| LINE | A text string, including spaces; no other parameters can follow this parameter |
| IFNAME | Interface name whose format varies depending on the platform; examples are: eth0 , Ethernet0 , ethernet0 , xe0 |
| A.B.C.D | IPv4 address |
| A.B.C.D/M | IPv4 address and mask/prefix |
| X:X::X:X | IPv6 address |
| X:X::X:X/M | IPv6 address and mask/prefix |
| HH:MM:SS | Time format |
| AA:NN | BGP community value |

Table 3. Variable placeholders (continued)

| Token | Description |
|--|---------------|
| XX:XX:XX:XX:XX:XX | MAC address |
| <1-5> <1-65535> <0-2147483647> <0-4294967295> | Numeric range |

Command Description Format

The [Table 4](#) table explains the sections used to describe each command in this reference.

Table 4. Command descriptions

| Section | Description |
|----------------|--|
| Command Name | The name of the command, followed by what the command does and when should it be used |
| Command Syntax | The syntax of the command |
| Parameters | Parameters and options for the command |
| Default | The state before the command is executed |
| Command Mode | The mode in which the command runs; see Command Modes (page 114) |
| Applicability | The command introduced in a specific release version and modified or updated in subsequent versions. |
| Example | An example of the command being executed |

Keyboard Operations

The [Table 5](#) table lists the operations you can perform from the keyboard.

Table 5. Keyboard operations

| Key combination | Operation |
|-----------------------|--|
| Left arrow or Ctrl+b | Moves one character to the left. When a command extends beyond a single line, you can press left arrow or Ctrl+b repeatedly to scroll toward the beginning of the line, or you can press Ctrl+a to go directly to the beginning of the line. |
| Right arrow or Ctrl-f | Moves one character to the right. When a command extends beyond a single line, you can press right arrow or Ctrl+f repeatedly to scroll toward the end of the line, or you can press Ctrl+e to go directly to the end of the line. |
| Esc, b | Moves back one word |
| Esc, f | Moves forward one word |
| Ctrl+e | Moves to end of the line |

Table 5. Keyboard operations (continued)

| Key combination | Operation |
|----------------------|--|
| Ctrl+a | Moves to the beginning of the line |
| Ctrl+u | Deletes the line |
| Ctrl+w | Deletes from the cursor to the previous whitespace |
| Alt+d | Deletes the current word |
| Ctrl+k | Deletes from the cursor to the end of line |
| Ctrl+y | Pastes text previously deleted with Ctrl+k, Alt+d, Ctrl+w, or Ctrl+u at the cursor |
| Ctrl+t | Transposes the current character with the previous character |
| Ctrl+c | Ignores the current line and redisplay the command prompt |
| Ctrl+z | Ends configuration mode and returns to exec mode |
| Ctrl+l | Clears the screen |
| Up Arrow or Ctrl+p | Scroll backward through command history |
| Down Arrow or Ctrl+n | Scroll forward through command history |

Show Command Modifiers



Note: The show command output included in the guides is for illustration purposes only. Based on the combination of features enabled and ongoing enhancements made to the commands, the output for these commands may vary. For instance, the actual command output may differ depending on the software version, configuration, and platform. Field names, values, and formats are subject to change.

You can use two tokens to modify the output of a **show** command. Enter a question mark to display these tokens:

```
# show users ?
| Output modifiers
> Output redirection
```

You can type the | (vertical bar character) to use output modifiers. For example:

```
> show rsvp | ?
begin      Begin with the line that matches
exclude    Exclude lines that match
include     Include lines that match
last       Last few lines
redirect   Redirect output
```

Begin Modifier

The **begin** modifier displays the output beginning with the first line that contains the input string (everything typed after the **begin** keyword). For example:

```
# show running-config | begin xe1
...skipping
interface xe1
ipv6 address fe80::204:75ff:fee6:5393/64
!
interface xe2
```

```

ipv6 address fe80::20d:56ff:fe96:725a/64
!
line con 0
login
!
end

```

You can specify a regular expression after the **begin** keyword. This example begins the output at a line with either “xe2” or “xe4”:

```

# show running-config | begin xe[2-4]

...skipping
interface xe2
 shutdown
!
interface xe4
 shutdown
!
interface svlan0.1
 no shutdown
!
route-map myroute permit 2
!
route-map mymap1 permit 10
!
route-map rmap1 permit 2
!
line con 0
 login
line vty 0 4
 login
!
end

```

Include Modifier

The **include** modifier includes only those lines of output that contain the input string. In the output below, all lines containing the word “input” are included:

```

# show interface xe1 | include input
  input packets 80434552, bytes 2147483647, dropped 0, multicast packets 0
  input errors 0, length 0, overrun 0, CRC 0, frame 0, fifo 1, missed 0

```

You can specify a regular expression after the **include** keyword. This examples includes all lines with “input” or “output”:

```

#show interface xe0 | include (in|out)put
  input packets 597058, bytes 338081476, dropped 0, multicast packets 0
  input errors 0, length 0, overrun 0, CRC 0, frame 0, fifo 0, missed 0
  output packets 613147, bytes 126055987, dropped 0
  output errors 0, aborted 0, carrier 0, fifo 0, heartbeat 0, window 0

```

Exclude Modifier

The **exclude** modifier excludes all lines of output that contain the input string. In the following output example, all lines containing the word “input” are excluded:

```

# show interface xe1 | exclude input
Interface xe1
 Scope: both
 Hardware is Ethernet, address is 0004.75e6.5393
 index 3 metric 1 mtu 1500 <UP,BROADCAST,RUNNING,MULTICAST>

```

```
VRF Binding: Not bound
Administrative Group(s): None
DSTE Bandwidth Constraint Mode is MAM
inet6 fe80::204:75ff:fee6:5393/64
  output packets 4438, bytes 394940, dropped 0
  output errors 0, aborted 0, carrier 0, fifo 0, heartbeat 0, window 0
  collisions 0
```

You can specify a regular expression after the **exclude** keyword. This example excludes lines with “output” or “input”:

```
show interface xe0 | exclude (in|out)put
Interface xe0
  Scope: both
  Hardware is Ethernet   Current HW addr: 001b.2139.6c4a
  Physical:001b.2139.6c4a Logical:(not set)
  index 2 metric 1 mtu 1500 duplex-full arp ageing timeout 3000
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Bandwidth 100m
  DHCP client is disabled.
  inet 10.1.2.173/24 broadcast 10.1.2.255
  VRRP Master of : VRRP is not configured on this interface.
  inet6 fe80::21b:21ff:fe39:6c4a/64
  collisions 0
```

Redirect Modifier

The **redirect** modifier writes the output into a file. The output is not displayed.

```
# show cli history | redirect /var/frame.txt
```

The output redirection token (>) does the same thing:

```
# show cli history >/var/frame.txt
```

Last Modifier

The **last** modifier displays the output of last few number of lines (As per the user input). The last number ranges from 1 to 9999.

For example:

```
#show running-config | last 10
```

String Parameters

The restrictions in [Table 6](#) apply for all string parameters used in OcNOS commands, unless some other restrictions are noted for a particular command.

Table 6. String parameter restrictions

| Restriction | Description |
|-------------------------------|--|
| Input length | 1965 characters or less |
| Restricted special characters | “?”, “,”, “>”, “ ”, and “=” The “ ” character is allowed only for the description command in interface mode. |

Command Modes

Commands are grouped into modes arranged in a hierarchy. Each mode has its own set of commands. The table below lists the command modes common to all protocols.


Table 7. Common Command Modes

| Name | Description |
|---------------------------|--|
| Execution mode | Also called <i>view</i> mode, this is the first mode to appear after you start the CLI. It is a base mode from where you can perform basic commands such as show, exit, quit, help, and enable. |
| Privileged execution mode | Also called <i>enable</i> mode, in this mode you can run additional basic commands such as debug, write, and show. |
| Configure mode | Also called <i>configure terminal</i> mode, in this mode you can run configuration commands and go into other modes such as interface, router, route map, key chain, and address family. Configure mode is single user. Only one user at a time can be in configure mode. |
| Interface mode | In this mode you can configure protocol-specific settings for a particular interface. Any setting you configure in this mode overrides a setting configured in router mode. |
| Router mode | This mode is used to configure router-specific settings for a protocol such as BGP or OSPF. |

Command Mode Tree

The diagram below shows the common command mode hierarchy.

Figure 1. Common command modes



To change modes:

1. Enter privileged executive mode by entering **enable** in Executive mode.
2. Enter configure mode by entering **configure terminal** in Privileged Executive mode.

The example below shows moving from executive mode to privileged executive mode to configure mode and finally to router mode:

```
> enable mypassword
# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
(config)# router ospf
(config-router)#
```



Note: Each protocol can have modes in addition to the common command modes. See the command reference for the respective protocol for details.

Transaction-based Command-line Interface

The OcNOS command line interface is transaction based:

- Any changes done in configure mode are stored in a separate *candidate* configuration that you can view with the `show transaction current` command.
- When a configuration is complete, apply the candidate configuration to the running configuration with the `commit` command.

- If a `commit` fails, no configuration is applied as the entire transaction is considered failed. You can continue to change the candidate configuration and then retry the `commit`.
- Discard the candidate configuration with the `abort transaction` command.
- Check the last aborted transaction with the `show transaction last-aborted` command.
- Multiple configurations cannot be removed with a single `.` You must remove each configuration followed by a `commit`.



Note: All commands MUST be executed only in the default CML shell (`cm1sh`). If you log in as root and start `imish`, then the system configurations will go out of sync. The `imish` shell is not supported and should not be started manually.

LABEL DISTRIBUTION PROTOCOL CONFIGURATION

| | |
|---|-----|
| LDP Configuration | 119 |
| Overview | 119 |
| Configure LDP | 120 |
| LDP MD5 Authentication | 126 |
| Configure LDP MD5 for Targeted LDP Session | 128 |
| Removing MD5 Authentication for LDP Session | 130 |
| LDP Session Protection | 133 |
| Configure Session Protection | 141 |
| LDP Inter-Area Configuration | 155 |
| Overview | 155 |
| Configure LDP Inter-Area | 155 |
| Validation | 161 |
| MPLS LDP-IGP Synchronization | 171 |
| Overview | 171 |
| Prerequisites | 171 |
| Topology | 171 |
| LDP-IGP Synchronization with OSPF | 172 |
| LDP-IGP Synchronization | 175 |
| LDP-IGP Synchronization with IS-IS | 177 |
| LDP-IGP SYNC Configuration | 182 |
| LDP-FRR Configuration | 186 |
| Topology | 187 |
| LDP-FRR with OSPF as IGP Configuration | 187 |
| LDP FRR with OSPF Configuration | 234 |
| LDP-FRR with ISIS as IGP Configuration | 278 |
| LDP FRR with IS-IS Configuration | 319 |
| Backup Path based on Route-Map Prefixes | 334 |
| TCP MSS configuration for LDP sessions | 351 |
| Overview | 351 |
| Prerequisites | 352 |
| Configuration | 352 |
| TCP MSS for LDP Session Command | 377 |
| Abbreviations | 378 |
| Glossary | 378 |
| LDP Remote Loop-Free Alternate (RLFA) | 379 |

| | |
|---|-----|
| RLFA Terminology | 379 |
| Establishing RLFA Tunnel | 380 |
| Configure LDP Remote Loop-Free Alternate (RLFA) | 381 |
| LDP MD5 Password For Auto Targeted Sessions | 416 |
| Configuration | 416 |
| Topology | 416 |
| Validation | 425 |
| LDP ECMP Configuration | 435 |
| Topology | 435 |
| Configurations | 435 |
| LDP Graceful Restart | 465 |
| Topology | 465 |
| Configuration | 465 |
| Validation | 467 |
| LDP Tunneling over RSVP-TE | 479 |
| Overview | 479 |
| Prerequisites | 479 |
| Limitations | 479 |
| Configuration for LDP Tunneling Over RSVP | 480 |
| Commands for LDP Tunneling over RSVP-TE | 490 |
| Glossary | 495 |
| Revert Timer for Revertive Switchover for LDP Pseudowires | 497 |
| Overview | 497 |
| Characteristics | 497 |
| Benefits | 497 |
| Prerequisites | 497 |
| Configuration | 497 |
| Revert Timer Commands | 512 |
| Implementation Examples | 512 |
| Troubleshooting Revert Timer for Revertive Switchover Issue | 512 |
| Glossary | 512 |

LDP Configuration

This chapter contains LDP (Label Distribution Protocol) configuration examples.

Overview

The Label Distribution Protocol (LDP) is a routing protocol used in MPLS technology. The LDP daemon (`ldpd`) uses NSM services to obtain routing information. Routers send Hello packets to establish Hello Adjacencies with other nearby routers. This opens the way for sessions between routers to be established during which routers exchange labels in preparation for forwarding packets.

LDP generates labels for and exchanges labels between peer routers. It works in with other routing protocols (RIP, OSPF and BGP) to create label-switched paths (LSP) used when forwarding packets. A label-switched path is the path taken by all packets that belong to the Forwarding Equivalence Class (FEC) corresponding to that LSP. This is analogous to establishing a virtual circuit in ATM (Asynchronous Transfer Mechanism). In this way, OcNOS LDP assigns labels to every destination address and destination prefix provided by OcNOS. The LDP interface to the MPLS forwarder adds labels to, and deletes labels from, the forwarding tables.

LDP Adjacencies

LDP defines a mechanism for discovering adjacent, LDP capable Label Switching Routers (LSR) that participate in label switching (adjacencies). Whenever a new router comes up it sends out a hello packet to a specified, multicast address announcing itself to the network. Every router directly connected to the network receives the packet. Receipt of a hello packet from another LSR creates a Hello Adjacency with that LSR. To create a Hello Adjacency with an LSR that cannot send/receive multicast packets, LDP allows a router to be manually configured to send unicast Hello packets to non-multicast LSRs. This non-multicast LSR is a targeted peer. Adjacencies are maintained by sending out periodic Hello packets to the multicast group and to all targeted peers. Hello packets are sent using UDP.

LDP Session

LDP capable LSRs establish a session before exchanging label information. All the session messages are sent using TCP to ensure reliable delivery. After the LSRs establish a session and negotiate options, a given pair of routers may exchange label information. The labels exchanged over a session are valid only during the lifetime of the session and routers release them when session is closed.

Forwarding Equivalence Class

A Forwarding Equivalence Class (FEC) section defines a set of packets that are forwarded on the same path by the MPLS network. Two common methods to define FEC are by advertising the IPv4 routes using:

- Host Address The LSR uses the address of the destination host to create this FEC. This means that all the packets going to this destination will take the same LSP.
- Prefix The LSR uses destination prefix to create this FEC. This means that all the packets take the LSP corresponding to the longest matching prefix.

Label Generation

An LDP Label is a 20-bit number the LSR uses to forward a packet to its destination. When an LSR creates a new FEC, the router generates new labels and distributes them to its peers. A router keeps both incoming and outgoing

labels in its database.

Label Distribution Modes

The OcNOS LDP implementation supports two label distribution modes:

- **Downstream Unsolicited** In this mode, next hop LSRs distribute labels to peers without waiting for a label request.
- **Downstream on Demand** In this mode, a LSR distributes a label to a peer only if there is a pending label request from the peer.

Label Retention Mode

The OcNOS LDP implementation supports two label retention modes:

- **Liberal Retention Mode** In this mode, the LSR retains all labels received from all sources. This mode helps in fast LSP setup in case of a change in next hop.
- **Conservative Retention Mode** In this mode, the LSR retains only those labels received from peers that are the next hop for a given FEC. This mode is used by LSRs that have a constraint on the number of labels that it can retain at any given time.

LSP Control

LSPs can be set up in the following two ways:

- **Ordered Control** In this mode, an LSR distributes a label for a FEC to its peer only if it has a corresponding label from its next hop or it is the egress node.
- **Independent Control** In this mode, an LSR may distribute a label to its peers without waiting for a corresponding label from its next hop.

Loop Detection

Loop detection can be enabled to detect routing loops in LSPs. There are two methods supported for the loop detection mechanism:

- **Hop Count** During setup of an LSP, the LSP passes hop count with the LSP setup messages. This hop count is incremented by each node router participating in LSP establishment. If the hop count exceeds the maximum configured value, the LSP setup process is stopped and a notification message is passed back to the message originator.
- **Path Vector** A path vector contains a list of LSR identifiers. This is passed as a part of LSP setup messages. Each LSR participating in the LSP establishment adds its own LSR identifier to the path vector. If an LSR finds its own identifier in the path vector, it drops the message and sends a message back to the originator.

The use of these messages ensures that a loop is detected while establishing a label switched path and before any data is passed over that LSP.

Configure LDP

The `enable-ldp ipv4` command is used to enable LDP for IPv4 on a specified interface, as follows:

- `enable-ldp ipv4` enables only IPv4 on the interface


For the examples covered in this section, the command `enable-ldp ipv4` is used.

Enable Label Switching

Running LDP on a system requires the following tasks:

1. Enabling label-switching on the interface on NSM.
2. Enabling LDP on an interface in the LDP daemon.
3. Running an IGP (Internal Gateway Protocol), for example, OSPF, to distribute reachability information within the MPLS cloud.
4. Configuring the transport address.


Figure 2. Basic LDP Topology



R1 - NSM

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.10.1/24 | Assign IP address to interface. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 192.168.0.1/32 secondary | Set the IP address of the loopback interface to 192.168.0.1/32. |
| (config-if)#commit | Commit the transaction. |

R1 - LDP

| | |
|--|---|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 192.168.0.1 | Set the router ID to IP address 192.168.0.1. |
| (config-router)#transport-address ipv4 192.168.0.1 | <p>Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface.</p> <div>  <p>Note: It is preferable to use the loopback address as transport address. In addition, use the parameter “ipv6” if you are configuring an IPv6 interface.</p> </div> |
| (config-router)#targeted-peer ipv4 192.168.0.3 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode. |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe1. |
| (config-if)#commit | Commit the transaction. |

R1 - OSPF

| | |
|---|---|
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.10.10.0/24 area 0 (config-router)#network 192.168.0.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

R2 - NSM

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 192.168.0.2/32 secondary | Set the IP address of the loopback interface to 192.168.0.2/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.10.2/24 | Assign IP address to interface |

| | |
|--------------------------------------|---|
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#ip address 10.10.11.1/24 | Assign IP address to interface. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#commit | Commit the transaction. |

R2 - LDP

| | |
|--|---|
| (config)#router ldp | Enter Router mode. |
| (config-router)#router-id 192.168.0.2 | Set the router ID to IP address 192.168.0.2. |
| (config-router)#transport-address ipv4 192.168.0.2 | <p>Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface.</p> <div data-bbox="889 787 961 850"> </div> <p>Note: Note: It is preferable to use the loopback address as transport address. In addition, use the parameter “ipv6” if you are configuring an IPv6 interface.</p> |
| (config-router)#exit | Exit Router mode and return to Configure mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface (xe1) . |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface (xe . |
| (config-if)#commit | Commit the transaction. |

R2 - OSPF


| | |
|---|---|
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.10.10.0/24 area 0 (config-router)#network 10.10.11.0/24 area 0 (config-router)#network 192.168.0.2/32 area 0 | Define the interfaces on which OSPF runs and associate the area ID (0) with them. |
| (config-router)#commit | Commit the transaction. |

R3 - NSM

| | |
|---------------------|-----------------------|
| #configure terminal | Enter configure mode. |
|---------------------|-----------------------|

| | |
|--|---|
| <code>(config)#interface lo</code> | Specify the loopback (lo) interface to be configured. |
| <code>(config-if)#ip address 192.168.0.3/32 secondary</code> | Set the IP address of the loopback interface to 192.168.0.3/32. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe2</code> | Specify the interface (xe2) to be configured. |
| <code>(config-if)#ip address 10.10.11.1/24</code> | Set the IP address of the interface to 10.10.11.1/24. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe2. |
| <code>(config-if)#commit</code> | Commit the transaction. |

R3 - LDP

| | |
|---|---|
| <code>(config)#router ldp</code> | Enter Router mode. |
| <code>(config-router)#router-id 192.168.0.3</code> | Set the router ID for IP address 192.168.0.3. |
| <code>(config-router)#transport-address ipv4 192.168.0.3</code> | <p>Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface.</p> <div data-bbox="891 940 964 1005">  </div> <p>Note: Note: It is preferable to use the loopback address as transport address. In addition, use the parameter “ipv6” if you are configuring an IPv6 interface.</p> |
| <code>(config-router)#targeted-peer ipv4 192.168.0.1</code> | Configure targeted peer. |
| <code>(config-router-targeted-peer)#exit</code> | Exit-targeted-peer-mode. |
| <code>(config-router)#exit</code> | Exit the Router mode and return to the Configure mode. |
| <code>(config)#interface xe2</code> | Enter interface mode. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP on xe2. |
| <code>(config-if)#commit</code> | Commit the transaction. |

R3 - OSPF

| | |
|---|---|
| <code>(config)#router ospf 100</code> | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| <code>(config-router)#network 10.10.11.0/24 area 0</code> <code>(config-router)#network 192.168.0.3/32 area 0</code> | Define the interfaces on which OSPF runs and associate the area ID (0) with them. |
| <code>(config-router)#commit</code> | Commit the transaction. |

Validation

PE1

```

PE1#show ldp session
Peer IP Address    IF Name  My Role  State          KeepAlive  UpTime
192.168.0.2        xe1      Passive  OPERATIONAL    30         00:18:59
192.168.0.3        xe1      Passive  OPERATIONAL    30         00:02:07

PE1#show ldp targeted-peer count
Num Targeted Peers: 1    [UP: 1]

PE1#show ldp session count
Multicast Peers   : 2    [UP: 1]
Targeted Peers    : 1    [UP: 0]
Total Sessions    : 2    [UP: 1]

PE1#show ldp routes
Prefix Addr  Nexthop Addr  Intf  Owner
10.10.10.0/24  0.0.0.0    xe1
10.10.11.0/24  10.10.10.2  xe1
192.168.0.1/32  0.0.0.0    lo
192.168.0.2/32  10.10.10.2  xe1
192.168.0.3/32  10.10.10.2  xe1

PE1#show ldp fec-ipv4 count
Num. IPv4 FEC(s): 5

```

P

```

P#show ldp session
Peer IP Address    IF Name  My Role  State          KeepAlive  UpTime
192.168.0.1        xe1      Active   OPERATIONAL    30         00:23:52
192.168.0.3        xe2      Passive  OPERATIONAL    30         00:23:49

P#show ldp session count
Multicast Peers   : 3    [UP: 2]
Targeted Peers    : 0    [UP: 0]
Total Sessions    : 2    [UP: 2]

P#show ldp routes
Prefix Addr  Nexthop Addr  Intf  Owner
10.10.10.0/24  0.0.0.0      xe1   connected
10.10.11.0/24  0.0.0.0      xe2   connected
192.168.0.1/32  10.10.10.1   xe1   ospf
192.168.0.2/32  0.0.0.0      lo    connected
192.168.0.3/32  10.10.11.2   xe2   ospf

```

PE2

```

PE2#show ldp session
Peer IP Address    IF Name  My Role  State          KeepAlive  UpTime
192.168.0.2        xe2      Active   OPERATIONAL    30         00:27:47
192.168.0.1        xe2      Active   OPERATIONAL    30         00:10:58

PE2#show ldp session count
Multicast Peers   : 2    [UP: 1] Targeted Peers    : 1    [UP: 1]
Total Sessions    : 2    [UP: 2]

OcNOS#show ldp targeted-peer count
Num Targeted Peers: 1    [UP: 1]

OcNOS#show ldp routes
Prefix Addr  Nexthop Addr  Intf  Owner

```

```
10.10.10.0/24 10.10.11.1 xe2 ospf
10.10.11.0/24 0.0.0.0 xe2 connected
192.168.0.1/32 10.10.11.1 xe17 ospf
192.168.0.2/32 10.10.11.1 xe2 ospf
192.168.0.3/32 0.0.0.0 lo connected
```

```
PE2#show ldp fec
fec fec-ipv4

PE2#show ldp fec-ipv4 count
-----
Num. IPv4 FEC(s): 5
-----
```


LDP MD5 Authentication

LDP MD5 configuration enables LDP MD5 password authentication on a per-peer basis.

Direct LDP Session

In this example, MD5 authentication is configured for a direct LDP session.

Figure 3. Topology for Direct Session MD5



R1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode. |
| (config)#router-id 1.1.1.1 | Configure the router id. |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#neighbor 10.10.10.2 auth md5 password 0 pwd1 | Configure the MD5 authentication and password, pwd1, for the neighbor, 10.10.10.2. |
| (config-router)#exit | Exit the Router mode and return to the |

| | |
|---|---|
| | Configure mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.10.1/24 | Set the IP address of the interface to 10.10.10.1/24.. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe1. |
| (config-if)#commit | Commit the transaction. |

R2

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode. |
| (config)#router-id 2.2.2.2 | Configure the router id. |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#neighbor 10.10.10.1 auth md5 password 0 pwd1 | Configure the MD5 authentication and password, pwd1, for the neighbor, 10.10.10.1. |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.10.2/24 | Set the IP address of the interface to 10.10.10.2/24. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe1. |
| (config-if)#commit | Commit the transaction. |

Validation**PE1**

```

PE1#show ldp session
Peer IP Address      IF Name   My Role   State      KeepAlive  UpTime
2.2.2.2              xe1       Passive   OPERATIONAL 30         00:14:53

PE1#show ldp session count
-----
Multicast Peers      : 1          [UP: 1]
Targeted Peers       : 1          [UP: 1]

```

Total Sessions : 1 [UP: 1]

PE2


```
PE2#show ldp session
Peer IP Address      IF Name  My Role  State      KeepAlive  UpTime
1.1.1.1              xe1      Active   OPERATIONAL 30         00:15:05

PE2#sh ldp session count
-----
Multicast Peers      : 1          [UP: 1]
Targeted Peers       : 1          [UP: 1]
Total Sessions       : 1          [UP: 1]
-----
```

Configure LDP MD5 for Targeted LDP Session

In this example, MD5 authentication is configured for the targeted LDP session established between R1 and R3.

Figure 4. Topology for Targeted Session MD5



PE1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode. |
| (config-router)#neighbor 10.10.10.2 auth md5 password 0 pwd1 | Configure the MD5 authentication and password, pwd1, for the neighbor, 10.10.10.2. |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configure the targeted peer IP address (R3 loopback address). |
| (config-router-targeted-peer)#exit | Exit targeted peer mode. |
| (config-router)#neighbor 3.3.3.3 auth md5 password 0 pwd2 | Configure the MD5 authentication and password, pwd2, for the targeted peer, 3.3.3.3. |
| (config-router)#exit | Exit the Router mode and return to the |

| | |
|---|--|
| | Configure mode. |
| (config)#interface lo | Specify the loopback interface (lo) to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/ 32. |
| (config-if)#commit | Commit the transaction. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.10.1/24 | Set the IP address of the interface to 10.10.10.1/24. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe1. |

P

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode to enable LDP. |
| (config)#router-id 2.2.2.2 | Configure the router ID. |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32 |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe2. |
| (config-if)#commit | Commit the transaction. |

PE2


| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode. |
| (config-router)#router-id 3.3.3.3 | Configure the router-id |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address to be used for a |

| | |
|--|--|
| | TCP session over which LDP will run on an IPv4. |
| <code>(config-router)#targeted-peer ipv4 1.1.1.1</code> | Configure the targeted peer IP address (R1 loopback address). |
| <code>(config-router-targeted-peer)#exit</code> | Exit targeted peer mode. |
| <code>(config-router)#neighbor 1.1.1.1 auth md5 password 0 pwd2</code> | Configure the MD5 authentication and password, pwd2, for the targeted peer, 1.1.1.1. |
| <code>(config-router)#exit</code> | Exit the Router mode and return to the Configure mode. |
| <code>(config)#interface lo</code> | Specify the loopback (lo) interface to be config. |
| <code>(config-if)#ip address 3.3.3.3/32 secondary</code> | Set the IP address of the loopback interface to 3.3.3.3/32 |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe2</code> | Specify the interface (xe2) to be configured. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe2. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP on interface xe2. |
| <code>(config-if)#commit</code> | Commit the transaction. |

Removing MD5 Authentication for LDP Session

This example shows removing the MD5 authentication configuration from an LDP session.

Figure 5. LDP Session Topology



PE1

| | |
|--|---|
| <code>#configure terminal</code> | Enter configure mode. |
| <code>(config)#router ldp</code> | Enter Router mode. |
| <code>(config-router)#no neighbor 10.10.10.2 auth md5</code> | Remove MD5 authentication for the neighbor, |

| | |
|------------------------|-------------------------|
| password | 10.10.10.2. |
| (config-router)#commit | Commit the transaction. |

PE2

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode. |
| (config-router)#no neighbor 10.10.10.1 auth md5 password | Remove MD5 authentication for the neighbor, 10.10.10.1. |
| (config-router)#commit | Commit the transaction. |
| (config)#interface lo | Specify the loopback (lo) interface to be config |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32 |
| (config-if)#exit | Exit lo interface |
| (config-if)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.10.2/24 | Set the IP address of the interface to 10.10.10.2/24 |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe1. |
| (config-if)#commit | Commit the transaction. |

Validation**LDP Session Count**

This example shows the number of configured LDP basic neighbors and targeted neighbors count.

```
#show ldp session count
-----
Basic sessions      : 100          [UP: 100]
Targeted sessions   : 500          [UP: 500]
Total Sessions      : 600          [UP: 600]
-----

#show ldp targeted-peer count
-----
Num Targeted Peers: 500          [UP: 500]
-----
```

FTN, SWAP, and POP Entries

This example shows forwarding table entries, SWAP entries and POP entries for IPV4 and IPV6 prefixes.

```
#show mpls forwarding-table count
-----
Num FTNs           : 300000        [UP: 3, INSTALLED: 300000]
Primary FTNs       : 300000        [UP: 3, INSTALLED: 300000]
Secondary FTNs     : 0             [UP: 0, INSTALLED: 0]
-----

Num IPV6 FTNs      : 300000        [UP: 300000, INSTALLED: 300000]
Primary IPV6 FTNs  : 300000        [UP: 300000, INSTALLED: 300000]
```

```

Secondary IPV6 FTNs : 0                      [UP: 0, INSTALLED: 0]
-----

#show mpls ilm-table count
-----
Num ILMs      : 300000                      [UP: 0, INSTALL: 300000]
Swap Entries  : 300000                      [UP: 0, INSTALL: 300000]
Pop Entries   : 0                          [UP: 0, INSTALL: 0]
VC Pop Entries : 0                          [UP: 0]
-----

```

MPLS LDP PING and TRACEROUTE

This example shows MPLS ping and trace route for LDP.

```

#show ip ospf neighbor
Total number of full neighbors: 1
OSPF process 0 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface      Instance ID
2.2.2.2        1    Full/DR         00:00:33    10.10.10.2     xe1            0

#show ldp session
Peer IP Address      IF Name    My Role    State        KeepAlive
2.2.2.2              xe1        Passive    OPERATIONAL  30

#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code   FEC                FTN-ID    Nhlfe-ID  Tunnel-id  Pri   LSP-Type    Out-Label    Out-
Intf   ELC   Nexthop
L>    2.2.2.2/32          1         1         -          -     LSP_
DEFAULT 3              xe1       No        10.10.10.2

#ping mpls ldp 2.2.2.2/32 detail
Sending 5 MPLS Echos to 2.2.2.2, timeout is 5 seconds
Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed
Type 'Ctrl+C' to abort
! seq_num = 1 2.2.2.2 1.73 ms
! seq_num = 2 2.2.2.2 1.46 ms
! seq_num = 3 2.2.2.2 0.64 ms
! seq_num = 4 2.2.2.20.65 ms
! seq_num = 5 2.2.2.20.62 ms
Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.62/1.18/1.73

#trace mpls ldp 2.2.2.2/32 detail
Tracing MPLS Label Switched Path to 2.2.2.2, timeout is 5 seconds
Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed
Type 'Ctrl+C' to abort
0 10.10.10.2 [Labels: 0]

```



```

! 1 2.2.2.2 0.69 ms

#ping mpls ldp 2.2.2.2/32 detail interval 5000 rep
reply-mode repeat
#ping mpls ldp 2.2.2.2/32 detail interval 5000 repeat 50
Sending 50 MPLS Echos to 2.2.2.2, timeout is 5 seconds
Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed
Type 'Ctrl+C' to abort
! seq_num = 1 2.2.2.2 0.70 ms
! seq_num = 2 2.2.2.2 0.73 ms
! seq_num = 3 2.2.2.2 0.71 ms
Success Rate is 100.00 percent (3/3)
round-trip min/avg/max = 0.70/0.71/0.73

```

LDP Session Protection

LDP Session Protection is an optimization feature. It is used when directly connected LDP peer sessions (via multicast) become unavailable but still have IP reachability over a different path. LDP bindings are kept in the LIB to save time from full synchronization when the direct connections comes back up.

There are two types of LDP connections:


- Direct LDP Session - directly connected LSR, one hop away.
- Targeted LDP Session - not directly connected LSR, multiple hops away.

By default if the directly connected LDP session loses connectivity to its peer, all bindings are flushed from the LIB. When interfaces come up and LDP sessions are re-established, LDP has to synchronize its label bindings.

Enable session protection for indirect link failures to protect the labels until the session is active when multicast adjacency gets deleted. When multicast and TLDP connection exists between the same peer, the interface down notification is received only on the router with the link failure but not in the peer. When the interface goes down the multicast adjacency and the labels get deleted without notifying the peer device. When the interface comes up before the adjacency hold timer expires in the peer node, multicast adjacency is added and advertises the labels to the peer but does not receive any labels from the peer as there is no change in the adjacency.

LDP Session Protection is an optimization, when enabled, will not flush the LIB when direct LDP sessions go down. As long as there exists another path to the LDP Peer, it will maintain the LIB synchronized using Targeted LDP Session. IGP will cause a reroute, but the label bindings will still be present from the old peer. When interfaces come back up, LDP will not need to synchronize since it maintains the state using the targeted sessions.


1. Running LDP Session Protection on a system requires the following tasks:
2. Enabling label-switching on the interface on NSM.
3. Enabling LDP on an interface in the LDP daemon.
4. Running an IGP (Internal Gateway Protocol), for example, OSPF, to distribute reachability information within the MPLS cloud.
5. Configuring the transport address.
6. Configuring LDP Session Protection.

Figure 6. Basic LDP Topology**PE1 - NSM**

| | |
|--------------------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe6 | Specify the interface (xe6) to be configured. |
| (config-if)#ip address 10.10.10.1/24 | Configure IPv4 address for xe6 |
| (config-if)#label-switching | Enable label switching on interface xe6. |
| (config)#interface xe12 | Specify the interface (xe12) to be configured. |
| (config-if)#ip address 30.30.30.1/24 | Configure IPv4 address for xe12 |
| (config-if)#label-switching | Enable label switching on interface xe12. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 1.1.1.1/32 | Set the IP address of the loopback interface to 1.1.1.1/32. |
| (config-if)#commit | Commit the transaction. |

PE1 - LDP

| | |
|--|--|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1 |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |

| | |
|-----------------------------|---|
| |  Note: Note: It is preferable to use the loopback address as transport address. In addition, use the parameter "ipv6" if you are configuring an IPv6 interface. |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#interface xe6 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe6. |
| (config)#interface xe12 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe12. |
| (config-if)#commit | Commit the transaction. |


PE1 - OSPF

| | |
|--|---|
| (config)#router ospf 1 | Configure the routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 1.1.1.1 | Configure Router ID |
| (config-router)#network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.10.10.1/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 30.30.30.1/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

P1 - NSM

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe47 | Specify the interface (xe47) to be configured. |
| (config-if)#ip address 10.10.10.2/24 | Configure IPv4 address for xe47 |
| (config-if)#label-switching | Enable label switching on interface xe47. |
| (config)#interface xe48 | Specify the interface (xe48) to be configured. |
| (config-if)#ip address 20.20.20.1/24 | Configure IPv4 address for xe48 |
| (config-if)#label-switching | Enable label switching on interface xe48. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| (config-if)#commit | Commit the transaction. |

P1 - LDP

| | |
|--|--|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. <div>  Note: Note: It is preferable to use the loopback address as transport address. In addition, use the parameter "ipv6" if you are configuring an IPv6 interface. </div> |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#interface xe47 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe47. |
| (config)#interface xe48 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe48. |
| (config-if)#commit | Commit the transaction. |

P1 - OSPF

| | |
|--|---|
| (config)#router ospf 1 | Configure the routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 2.2.2.2 | Configure Router ID |
| (config-router)#network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.10.10.2/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 20.20.20.1/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

P2 - NSM

| | |
|--------------------------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#interface xe48 | Specify the interface (xe48) to be configured. |
| (config-if)#ip address 30.30.30.2/24 | Configure IPv4 address for xe48 |
| (config-if)#label-switching | Enable label switching on interface xe48. |
| (config)#interface xe5 | Specify the interface (xe5) to be configured. |

| | |
|---|---|
| (config-if)#ip address 40.40.40.1/24 | Configure IPv4 address for xe5 |
| (config-if)#label-switching | Enable label switching on interface xe5. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 4.4.4.4/32 secondary | Set the IP address of the loopback interface to 4.4.4.4/32. |
| (config-if)#commit | Commit the transaction. |

P2 - LDP

| | |
|--|--|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 4.4.4.4 | Set the router ID to IP address 4.4.4.4 |
| (config-router)#transport-address ipv4 4.4.4.4 | <p>Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface.</p> <div data-bbox="867 781 943 848" data-label="Image"> </div> <p>Note: Note: It is preferable to use the loopback address as transport address. In addition, use the parameter "ipv6" if you are configuring an IPv6 interface.</p> |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#interface xe48 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe48. |
| (config)#interface xe5 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe5. |
| (config-if)#commit | Commit the transaction. |

P2 - OSPF

| | |
|--|---|
| (config)#router ospf 1 | Configure the routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 4.4.4.4 | Configure Router ID |
| (config-router)#network 4.4.4.4/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 20.20.20.2/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 30.30.30.1/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

PE2 - NSM

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe10 | Specify the interface (xe10) to be configured. |
| (config-if)#ip address 20.20.20.2/24 | Configure IPv4 address for xe10 |
| (config-if)#label-switching | Enable label switching on interface xe10. |
| (config)#interface xe5 | Specify the interface (xe5) to be configured. |
| (config-if)#ip address 40.40.40.2/24 | Configure IPv4 address for xe5 |
| (config-if)#label-switching | Enable label switching on interface xe5. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 3.3.3.3/32 secondary | Set the IP address of the loopback interface to 3.3.3.3/32. |
| (config-if)#commit | Commit the transaction. |

PE2 - LDP

| | |
|--|---|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 3.3.3.3 | Set the router ID to IP address 3.3.3.3 |
| (config-router)#transport-address ipv4 3.3.3.3 | <p>Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface.</p> <div data-bbox="867 1087 943 1155" data-label="Image"></div> <p>Note: Note: It is preferable to use the loopback address as transport address. In addition, use the parameter "ipv6" if you are configuring an IPv6 interface.</p> |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#interface xe10 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe10. |
| (config)#interface xe5 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe5. |
| (config-if)#commit | Commit the transaction. |

PE2 - OSPF

| | |
|---|---|
| (config)#router ospf 1 | Configure the routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 3.3.3.3 | Configure Router ID |
| (config-router)#network 3.3.3.3/32 area 0 | Define the interface on which OSPF runs and |

| | |
|--|---|
| | associate the area ID (0) with the interface. |
| (config-router)#network 20.20.20.2/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 40.40.40.2/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

Validation

Without session protection enabled

Verify that session protection status is not shown when session protection not enabled.

```

PE1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C       1.1.1.1/32 is directly connected, lo, 00:04:22
O       2.2.2.2/32 [110/2] via 10.10.10.2, xe12, 00:03:03
O       3.3.3.3/32 [110/3] via 10.10.10.2, xe12, 00:02:49
O       4.4.4.4/32 [110/31] via 30.30.30.2, xe6, 00:02:17
C       10.10.10.0/24 is directly connected, xe12, 00:03:48
O       20.20.20.0/24 [110/2] via 10.10.10.2, xe12, 00:03:03
C       30.30.30.0/24 is directly connected, xe6, 00:03:02
O       40.40.40.0/24 [110/31] via 30.30.30.2, xe6, 00:02:17
C       127.0.0.0/8 is directly connected, lo, 00:04:22

Gateway of last resort is not set

PE1#show ldp session
Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
4.4.4.4              xe6        Passive    OPERATIONAL 30         00:02:25
2.2.2.2              xe12       Passive    OPERATIONAL 30         00:03:11
PE1#show ldp targeted-peers
PE1#show ldp session 2.2.2.2
Session state        : OPERATIONAL
Session role         : Passive
TCP Connection       : Established
IP Address for TCP   : 2.2.2.2
Interface being used : xe12
Peer LDP ID          : 2.2.2.2:0
Peer LDP Password    : Not Set
Adjacencies          : 10.10.10.2
Advertisement mode    : Downstream Unsolicited
Label retention mode  : Liberal
Graceful Restart     : Not Capable
Keepalive Timeout    : 30
Reconnect Interval   : 15
Address List received : 2.2.2.2
                     10.10.10.2
                     20.20.20.1
                     254.128.0.0
Received Labels :      Fec          Label          Maps To
                  IPV4:3.3.3.3/32    52480          24963
                  IPV4:20.20.20.0/24 impl-null       24964

```

```

                IPV4:10.10.10.0/24      impl-null    none
                IPV4:2.2.2.2/32         impl-null    24962
Sent Labels :   Fec                     Label         Maps To
                IPV4:40.40.40.0/24      24961         impl-null
                IPV4:4.4.4.4/32         24960         impl-null
                IPV4:30.30.30.0/24      impl-null    none
                IPV4:10.10.10.0/24      impl-null    none
                IPV4:1.1.1.1/32         impl-null    none

```

PE1#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|----------|----------|------------|-----|----------|-----------|------|
| Intf | ELC | Nextthop | | | | | | |
| L> | 2.2.2.2/32 | 1 | 2 | - | Yes | LSP_ | | |
| DEFAULT | 3 | xe12 | No | 10.10.10.2 | | | | |
| L> | 3.3.3.3/32 | 3 | 5 | - | Yes | LSP_ | | |
| DEFAULT | 52480 | xe12 | No | 10.10.10.2 | | | | |
| L> | 4.4.4.4/32 | 4 | 7 | - | Yes | LSP_ | | |
| DEFAULT | 3 | xe6 | No | 30.30.30.2 | | | | |
| L> | 20.20.20.0/24 | 2 | 3 | - | Yes | LSP_ | | |
| DEFAULT | 3 | xe12 | No | 10.10.10.2 | | | | |
| L> | 40.40.40.0/24 | 5 | 8 | - | Yes | LSP_ | | |
| DEFAULT | 3 | xe6 | No | 30.30.30.2 | | | | |

PE1#show mpls ftn-table

Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe12, out label: 3
 Nextthop addr: 10.10.10.2 cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 3, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 4
 Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: xe12, out label: 52480
 Nextthop addr: 10.10.10.2 cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 4, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 6, owner: N/A, Stale: NO, out intf: xe6, out label: 3
 Nextthop addr: 30.30.30.2 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 20.20.20.0/24, id: 2, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe12, out label: 3
 Nextthop addr: 10.10.10.2 cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 40.40.40.0/24, id: 5, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0


```
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 6, owner: N/A, Stale: NO, out intf: xe6, out label: 3
Nexthop addr: 30.30.30.2      cross connect ix: 4, op code: Push
```

```
PE1#show mpls ilm-table
```

```
Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM, T - MPLS-TP, s - Stitched ILM
        S - SNMP, L - LDP, R - RSVP, C - CRLDP
        B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
        O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
        P - SR Policy, U - unknown
```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|--------|----------|-----------|---------|-----------------|
| Intf/VRF | Nexthop | | LSP-Type | | | |
| L> | 3.3.3.3/32 | 4 | 24963 | 52480 | N/A | xe12 10.10.10.2 |
| | LSP_DEFAULT | | | | | |
| L> | 40.40.40.0/24 | 2 | 24961 | 3 | N/A | xe6 30.30.30.2 |
| | LSP_DEFAULT | | | | | |
| L> | 4.4.4.4/32 | 1 | 24960 | 3 | N/A | xe6 30.30.30.2 |
| | LSP_DEFAULT | | | | | |
| L> | 2.2.2.2/32 | 3 | 24962 | 3 | N/A | xe12 10.10.10.2 |
| | LSP_DEFAULT | | | | | |
| L> | 20.20.20.0/24 | 5 | 24964 | 3 | N/A | xe12 10.10.10.2 |
| | LSP_DEFAULT | | | | | |

```
PE1#show ldp fec
```

```
LSR codes : E/N - LSR is egress/non-egress for this FEC,
            L - LSR received a label for this FEC,
            > - LSR will use this route for the FEC
```

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|---------------|------|--------------|-----------|-----|--------------|
| 1.1.1.1/32 | E > | non-existent | none | No | connected |
| 2.2.2.2/32 | NL> | 2.2.2.2 | impl-null | No | 10.10.10.2 |
| 3.3.3.3/32 | NL | 4.4.4.4 | 24325 | No | no nexthop |
| | NL> | 2.2.2.2 | 52480 | No | 10.10.10.2 |
| 4.4.4.4/32 | NL> | 4.4.4.4 | impl-null | No | 30.30.30.2 |
| 10.10.10.0/24 | NL | 2.2.2.2 | impl-null | No | connected |
| | E > | non-existent | none | No | connected |
| 20.20.20.0/24 | NL | 4.4.4.4 | 24326 | No | no nexthop |
| | NL> | 2.2.2.2 | impl-null | No | 10.10.10.2 |
| 30.30.30.0/24 | NL | 4.4.4.4 | impl-null | No | connected |
| | E > | non-existent | none | No | connected |
| 40.40.40.0/24 | NL> | 4.4.4.4 | impl-null | No | 30.30.30.2 |

Configure Session Protection



Note: Recommended to configure both ends.

Configure session protection under LDP in both nodes.

PE1

| | |
|------------------------------------|--|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#session-protection | Session-protection protect label indefinitely if no timer mentioned. |
| (config-router)#commit | Commit and exit |

P1

| | |
|------------------------------------|--|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#session-protection | Session-protection protect label indefinitely if no timer mentioned. |
| (config-router)#commit | Commit and exit |

Validation**With session protection command enabled**

Verify that session protection status shown once session protection enabled in both peer nodes.

```

PE1#show ldp session
Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
4.4.4.4              xe6        Passive    OPERATIONAL 30      00:05:46
2.2.2.2              xe12       Passive    OPERATIONAL 30      00:06:32
PE1#show ldp targeted-peers
IP Address           Interface
2.2.2.2              xe12
4.4.4.4              xe6
PE1#show ldp session 2.2.2.2
Session state        : OPERATIONAL
Session role         : Passive
TCP Connection       : Established
IP Address for TCP   : 2.2.2.2
Interface being used : xe12
Peer LDP ID          : 2.2.2.2:0
Peer LDP Password    : Not Set
Adjacencies          : 10.10.10.2
                     : 2.2.2.2
Advertisement mode    : Downstream Unsolicited
Label retention mode  : Liberal
Graceful Restart     : Not Capable
Keepalive Timeout    : 30
Reconnect Interval   : 15
Session protection   : Ready
Address List received : 2.2.2.2
                     : 10.10.10.2
                     : 20.20.20.1
                     : 254.128.0.0
Received Labels :    Fec          Label          Maps To
                   IPV4:3.3.3.3/32    52480          24963
                   IPV4:20.20.20.0/24 impl-null      24964
                   IPV4:10.10.10.0/24 impl-null      none
                   IPV4:2.2.2.2/32    impl-null      24962
Sent Labels :      Fec          Label          Maps To
                   IPV4:40.40.40.0/24 24961          impl-null
                   IPV4:4.4.4.4/32    24960          impl-null
                   IPV4:30.30.30.0/24 impl-null      none
                   IPV4:10.10.10.0/24 impl-null      none
                   IPV4:1.1.1.1/32    impl-null      none
PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport

Code  FEC          Nexthop      FTN-ID  Nhlfe-ID  Tunnel-id  Pri  LSP-Type  Out-Label  Out-
Intf  ELC
L>    2.2.2.2/32    xe12       1        2        -        Yes  LSP_
DEFAULT 3          3          No       10.10.10.2
L>    3.3.3.3/32    3          5        -        Yes  LSP_

```

```

DEFAULT 52480 xe12 No 10.10.10.2
L> 4.4.4.4/32 4 7 - Yes LSP_
DEFAULT 3 xe6 No 30.30.30.2
L> 20.20.20.0/24 2 3 - Yes LSP_
DEFAULT 3 xe12 No 10.10.10.2
L> 40.40.40.0/24 5 8 - Yes LSP_
DEFAULT 3 xe6 No 30.30.30.2

```

PE1#show mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM
 K - CLI ILM, T - MPLS-TP, s - Stitched ILM
 S - SNMP, L - LDP, R - RSVP, C - CRLDP
 B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
 O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
 P - SR Policy, U - unknown

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|--------|----------|-----------|---------|-----------------|
| Intf/VRF | NextHop | | LSP-Type | | | |
| L> | 3.3.3.3/32 | 4 | 24963 | 52480 | N/A | xe12 10.10.10.2 |
| | LSP_DEFAULT | | | | | |
| L> | 40.40.40.0/24 | 2 | 24961 | 3 | N/A | xe6 30.30.30.2 |
| | LSP_DEFAULT | | | | | |
| L> | 4.4.4.4/32 | 1 | 24960 | 3 | N/A | xe6 30.30.30.2 |
| | LSP_DEFAULT | | | | | |
| L> | 2.2.2.2/32 | 3 | 24962 | 3 | N/A | xe12 10.10.10.2 |
| | LSP_DEFAULT | | | | | |
| L> | 20.20.20.0/24 | 5 | 24964 | 3 | N/A | xe12 10.10.10.2 |
| | LSP_DEFAULT | | | | | |

PE1#show mpls ftn-table

Primary FTM entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe12, out label: 3
 Nexthop addr: 10.10.10.2 cross connect ix: 1, op code: Push

Primary FTM entry with FEC: 3.3.3.3/32, id: 3, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 4
 Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: xe12, out label: 52480
 Nexthop addr: 10.10.10.2 cross connect ix: 2, op code: Push

Primary FTM entry with FEC: 4.4.4.4/32, id: 4, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 6, owner: N/A, Stale: NO, out intf: xe6, out label: 3
 Nexthop addr: 30.30.30.2 cross connect ix: 4, op code: Push

Primary FTM entry with FEC: 20.20.20.0/24, id: 2, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe12, out label: 3
 Nexthop addr: 10.10.10.2 cross connect ix: 1, op code: Push

Primary FTM entry with FEC: 40.40.40.0/24, id: 5, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 6, owner: N/A, Stale: NO, out intf: xe6, out label: 3
 Nexthop addr: 30.30.30.2 cross connect ix: 4, op code: Push

PE1#show ldp fec

LSR codes : E/N - LSR is egress/non-egress for this FEC,
 L - LSR received a label for this FEC,
 > - LSR will use this route for the FEC

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|---------------|------|--------------|-----------|-----|--------------|
| 1.1.1.1/32 | E > | non-existent | none | No | connected |
| 2.2.2.2/32 | NL> | 2.2.2.2 | impl-null | No | 10.10.10.2 |
| 3.3.3.3/32 | NL | 4.4.4.4 | 24325 | No | no nexthop |
| | NL> | 2.2.2.2 | 52480 | No | 10.10.10.2 |
| 4.4.4.4/32 | NL> | 4.4.4.4 | impl-null | No | 30.30.30.2 |
| 10.10.10.0/24 | NL | 2.2.2.2 | impl-null | No | connected |
| | E > | non-existent | none | No | connected |
| 20.20.20.0/24 | NL | 4.4.4.4 | 24326 | No | no nexthop |
| | NL> | 2.2.2.2 | impl-null | No | 10.10.10.2 |
| 30.30.30.0/24 | NL | 4.4.4.4 | impl-null | No | connected |
| | E > | non-existent | none | No | connected |
| 40.40.40.0/24 | NL> | 4.4.4.4 | impl-null | No | 30.30.30.2 |

P1#show ldp session

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 3.3.3.3 | xe5 | Passive | OPERATIONAL | 30 | 00:05:40 |
| 1.1.1.1 | xe48 | Active | OPERATIONAL | 30 | 00:06:43 |

P1#show ldp targeted-peers

| IP Address | Interface |
|------------|-----------|
| 1.1.1.1 | xe48 |
| 3.3.3.3 | xe5 |

P1#show ldp session 1.1.1.1

Session state : OPERATIONAL
 Session role : Active
 TCP Connection : Established
 IP Address for TCP : 1.1.1.1
 Interface being used : xe48
 Peer LDP ID : 1.1.1.1:0
 Peer LDP Password : Not Set
 Adjacencies : 10.10.10.1
 1.1.1.1
 Advertisement mode : Downstream Unsolicited
 Label retention mode : Liberal
 Graceful Restart : Not Capable
 Keepalive Timeout : 30
 Reconnect Interval : 15
 Session protection : Ready
 Address List received : 1.1.1.1
 10.10.10.1
 30.30.30.1
 254.128.0.0

| Received Labels : | Fec | Label | Maps To |
|-------------------|--------------------|-----------|-----------|
| | IPV4:4.4.4.4/32 | 24960 | 52482 |
| | IPV4:40.40.40.0/24 | 24961 | 52484 |
| | IPV4:30.30.30.0/24 | impl-null | 52483 |
| | IPV4:10.10.10.0/24 | impl-null | none |
| | IPV4:1.1.1.1/32 | impl-null | 52481 |
| Sent Labels : | Fec | Label | Maps To |
| | IPV4:3.3.3.3/32 | 52480 | impl-null |
| | IPV4:20.20.20.0/24 | impl-null | none |
| | IPV4:10.10.10.0/24 | impl-null | none |
| | IPV4:2.2.2.2/32 | impl-null | none |

P1#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,

U - unknown FTN

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|--------|------------|-----------|-----|----------|-----------|------|
| Intf | ELC Nexthop | | | | | | | |
| L> | 1.1.1.1/32 | 1 | 2 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe48 | No | 10.10.10.1 | | | | | |
| L> | 3.3.3.3/32 | 5 | 9 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe5 | No | 20.20.20.2 | | | | | |
| L> | 4.4.4.4/32 | 3 | 5 | - | Yes | LSP_ | | |
| DEFAULT | 24960 xe48 | No | 10.10.10.1 | | | | | |
| L> | 30.30.30.0/24 | 2 | 3 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe48 | No | 10.10.10.1 | | | | | |
| L> | 40.40.40.0/24 | 4 | 7 | - | Yes | LSP_ | | |
| DEFAULT | 24961 xe48 | No | 10.10.10.1 | | | | | |

Pl#show mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM
 K - CLI ILM, T - MPLS-TP, s - Stitched ILM
 S - SNMP, L - LDP, R - RSVP, C - CRLDP
 B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
 U - unknown

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|--------|----------|-----------|---------|-----------------|
| Intf/VRF | Nexthop | | LSP-Type | | | |
| L> | 30.30.30.0/24 | 4 | 52483 | 3 | N/A | xe48 10.10.10.1 |
| | LSP_DEFAULT | | | | | |
| L> | 1.1.1.1/32 | 2 | 52481 | 3 | N/A | xe48 10.10.10.1 |
| | LSP_DEFAULT | | | | | |
| L> | 3.3.3.3/32 | 1 | 52480 | 3 | N/A | xe5 20.20.20.2 |
| | LSP_DEFAULT | | | | | |
| L> | 4.4.4.4/32 | 3 | 52482 | 24960 | N/A | xe48 10.10.10.1 |
| | LSP_DEFAULT | | | | | |
| L> | 40.40.40.0/24 | 5 | 52484 | 24961 | N/A | xe48 10.10.10.1 |
| | LSP_DEFAULT | | | | | |

Pl#show mpls ftn-table

Primary FTN entry with FEC: 1.1.1.1/32, id: 1, row status: Active
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A
 Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe48, out label: 3
 Nexthop addr: 10.10.10.1 cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 5, row status: Active
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A
 Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 8
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 8, owner: N/A, Stale: NO, out intf: xe5, out label: 3
 Nexthop addr: 20.20.20.2 cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 3, row status: Active
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A
 Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4
 Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: xe48, out label: 24960
 Nexthop addr: 10.10.10.1 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 30.30.30.0/24, id: 2, row status: Active
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A
 Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe48, out label: 3

```
Nexthop addr: 10.10.10.1      cross connect ix: 1, op code: Push
```

```
Primary FTN entry with FEC: 40.40.40.0/24, id: 4, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: xe48, out label: 24961
Nexthop addr: 10.10.10.1      cross connect ix: 4, op code: Push
```

```
Pl#show ldp fec
```

```
LSR codes      : E/N - LSR is egress/non-egress for this FEC,
                  L - LSR received a label for this FEC,
                  > - LSR will use this route for the FEC
```

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|---------------|------|--------------|-----------|-----|--------------|
| 1.1.1.1/32 | NL> | 1.1.1.1 | impl-null | No | 10.10.10.1 |
| 2.2.2.2/32 | E > | non-existent | none | No | connected |
| 3.3.3.3/32 | NL> | 3.3.3.3 | impl-null | No | 20.20.20.2 |
| 4.4.4.4/32 | NL> | 1.1.1.1 | 24960 | No | 10.10.10.1 |
| 10.10.10.0/24 | NL | 1.1.1.1 | impl-null | No | connected |
| | E > | non-existent | none | No | connected |
| 20.20.20.0/24 | NL | 3.3.3.3 | impl-null | No | connected |
| | E > | non-existent | none | No | connected |
| 30.30.30.0/24 | NL> | 1.1.1.1 | impl-null | No | 10.10.10.1 |
| 40.40.40.0/24 | NL | 3.3.3.3 | impl-null | No | no nexthop |
| | NL> | 1.1.1.1 | 24961 | No | 10.10.10.1 |

Perform Link failure and check labels are retained until peer is reachable through alternate path.

| | |
|-------------------------|-----------------------|
| (config)#interface xe12 | Enter interface mode. |
| (config-if)#shutdown | Shutdown the link. |
| (config)#commit | commit. |

Link Down Validation

After link down

```
PE1#show ip route
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default
```

```
IP Route Table for VRF "default"
```

```
C      1.1.1.1/32 is directly connected, lo, 00:14:17
O      2.2.2.2/32 [110/33] via 30.30.30.2, xe6, 00:03:38
O      3.3.3.3/32 [110/32] via 30.30.30.2, xe6, 00:03:38
O      4.4.4.4/32 [110/31] via 30.30.30.2, xe6, 00:12:12
O      20.20.20.0/24 [110/32] via 30.30.30.2, xe6, 00:03:38
C      30.30.30.0/24 is directly connected, xe6, 00:12:57
O      40.40.40.0/24 [110/31] via 30.30.30.2, xe6, 00:12:12
C      127.0.0.0/8 is directly connected, lo, 00:14:17
```

```
PE1#show ldp session
```

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 4.4.4.4 | xe6 | Passive | OPERATIONAL | 30 | 00:10:10 |
| 2.2.2.2 | xe6 | Passive | OPERATIONAL | 30 | 00:10:56 |

```
PE1#show ldp targeted-peers
```

```

IP Address      Interface
2.2.2.2         xe6
4.4.4.4         xe6/
PE1#show ldp session 2.2.2.2
Session state    : OPERATIONAL
Session role     : Passive
TCP Connection   : Established
IP Address for TCP : 2.2.2.2
Interface being used : xe6
Peer LDP ID      : 2.2.2.2:0
Peer LDP Password : Not Set
Adjacencies      : 2.2.2.2
Advertisement mode : Downstream Unsolicited
Label retention mode : Liberal
Graceful Restart : Not Capable
Keepalive Timeout : 30
Reconnect Interval : 15
Session protection : Protecting
Address List received : 2.2.2.2
                    20.20.20.1
                    254.128.0.0

Received Labels :      Fec          Label      Maps To
                    IPV4:3.3.3.3/32    52480      none
                    IPV4:20.20.20.0/24 impl-null  none
                    IPV4:10.10.10.0/24 impl-null  none
                    IPV4:2.2.2.2/32    impl-null  none
Sent Labels :      Fec          Label      Maps To
                    IPV4:40.40.40.0/24 24961      impl-null
                    IPV4:4.4.4.4/32    24960      impl-null
                    IPV4:30.30.30.0/24 impl-null  none
                    IPV4:10.10.10.0/24 impl-null  none
                    IPV4:1.1.1.1/32    impl-null  none

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
        B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport

```

| Code | FEC | FTN-ID | Nhlife-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|---------|-----------|------------|-----|----------|-----------|------|
| Intf | ELC | Nexthop | | | | | | |
| L> | 2.2.2.2/32 | 3 | 9 | - | Yes | LSP_ | | |
| DEFAULT | 24321 | xe6 | No | 30.30.30.2 | | | | |
| L> | 3.3.3.3/32 | 1 | 2 | - | Yes | LSP_ | | |
| DEFAULT | 24325 | xe6 | No | 30.30.30.2 | | | | |
| L> | 4.4.4.4/32 | 4 | 7 | - | Yes | LSP_ | | |
| DEFAULT | 3 | xe6 | No | 30.30.30.2 | | | | |
| L> | 20.20.20.0/24 | 2 | 4 | - | Yes | LSP_ | | |
| DEFAULT | 24326 | xe6 | No | 30.30.30.2 | | | | |
| L> | 40.40.40.0/24 | 5 | 8 | - | Yes | LSP_ | | |
| DEFAULT | 3 | xe6 | No | 30.30.30.2 | | | | |

```

PE1#show mpls ftn-table
Primary FTN entry with FEC: 2.2.2.2/32, id: 3, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 5
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 5, owner: LDP, Stale: NO, out intf: xe6, out label: 24321
Nexthop addr: 30.30.30.2      cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 1, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

```

```

    Out-segment with ix: 1, owner: LDP, Stale: NO, out intf: xe6, out label: 24325
    Nexthop addr: 30.30.30.2          cross connect ix: 1, op code: Push

```

```

Primary FTN entry with FEC: 4.4.4.4/32, id: 4, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 6, owner: N/A, Stale: NO, out intf: xe6, out label: 3
    Nexthop addr: 30.30.30.2          cross connect ix: 4, op code: Push

```

```

Primary FTN entry with FEC: 20.20.20.0/24, id: 2, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 3
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 3, owner: LDP, Stale: NO, out intf: xe6, out label: 24326
    Nexthop addr: 30.30.30.2          cross connect ix: 2, op code: Push

```

```

Primary FTN entry with FEC: 40.40.40.0/24, id: 5, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 6, owner: N/A, Stale: NO, out intf: xe6, out label: 3
    Nexthop addr: 30.30.30.2          cross connect ix: 4, op code: Push

```

```
PE1#show mpls ilm-table
```

```
Codes: > - installed ILM, * - selected ILM, p - stale ILM
```

```

  K - CLI ILM, T - MPLS-TP, s - Stitched ILM
  S - SNMP, L - LDP, R - RSVP, C - CRLDP
  B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
  O - OSPF/OSPF6 SR, i - ISIS_SR, k - SR CLI
  P - SR Policy, U - unknown

```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|--------|----------|-----------|---------|----------------|
| Intf/VRF | Nexthop | | LSP-Type | | | |
| L> | 40.40.40.0/24 | 2 | 24961 | 3 | N/A | xe6 30.30.30.2 |
| | LSP_DEFAULT | | | | | |
| L> | 4.4.4.4/32 | 1 | 24960 | 3 | N/A | xe6 30.30.30.2 |
| | LSP_DEFAULT | | | | | |

```
P1#show ldp session
```

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 3.3.3.3 | xe5 | Passive | OPERATIONAL | 30 | 00:11:12 |
| 1.1.1.1 | xe5 | Active | OPERATIONAL | 30 | 00:12:15 |

```
P1#show ldp targeted-peers
```

| IP Address | Interface |
|------------|-----------|
| 1.1.1.1 | xe5 |
| 3.3.3.3 | xe5 |

```
P1#show ldp session 1.1.1.1
```

```

Session state      : OPERATIONAL
Session role       : Active
TCP Connection     : Established
IP Address for TCP : 1.1.1.1
Interface being used : xe5
Peer LDP ID        : 1.1.1.1:0
Peer LDP Password   : Not Set
Adjacencies        : 1.1.1.1
Advertisement mode   : Downstream Unsolicited
Label retention mode : Liberal
Graceful Restart    : Not Capable
Keepalive Timeout   : 30

```



```

Reconnect Interval      : 15
Session protection      : Protecting
Address List received   : 1.1.1.1
                        30.30.30.1
                        254.128.0.0

```

```

Received Labels :      Fec          Label          Maps To
                    IPV4:4.4.4.4/32    24960          52482
                    IPV4:40.40.40.0/24 24961          52484
                    IPV4:30.30.30.0/24 impl-null      52483
                    IPV4:10.10.10.0/24 impl-null      none
                    IPV4:1.1.1.1/32    impl-null      52481
Sent Labels :      Fec          Label          Maps To
                    IPV4:3.3.3.3/32    52480          impl-null
                    IPV4:20.20.20.0/24 impl-null      none
                    IPV4:10.10.10.0/24 impl-null      none
                    IPV4:2.2.2.2/32    impl-null      none

```

Pl#show mpls forwarding-table

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
        B - BGP FTN, K - CLI FTN, t - tunnel,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN

```

| Code | FEC | Nexthop | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|---------|--------|------------|-----------|-----|----------|-----------|------|
| L> | 1.1.1.1/32 | | 2 | 3 | - | Yes | LSP_ | | |
| DEFAULT | 24965 | xe5 | No | 20.20.20.2 | | | | | |
| L> | 3.3.3.3/32 | | 5 | 9 | - | Yes | LSP_ | | |
| DEFAULT | 3 | xe5 | No | 20.20.20.2 | | | | | |
| L> | 4.4.4.4/32 | | 3 | 5 | - | Yes | LSP_ | | |
| DEFAULT | 24966 | xe5 | No | 20.20.20.2 | | | | | |
| L> | 30.30.30.0/24 | | 4 | 7 | - | Yes | LSP_ | | |
| DEFAULT | 24967 | xe5 | No | 20.20.20.2 | | | | | |
| L> | 40.40.40.0/24 | | 1 | 1 | - | Yes | LSP_ | | |
| DEFAULT | 3 | xe5 | No | 20.20.20.2 | | | | | |

Pl#show mpls ilm-table

```

Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM, T - MPLS-TP, s - Stitched ILM
        S - SNMP, L - LDP, R - RSVP, C - CRLDP
        B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
        U - unknown

```

| Code | FEC/VRF | L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|------|---------------|-------------|--------|----------|-----------|---------|------------|
| L> | 4.4.4.4/32 | | 3 | 52482 | Nolabel | N/A | N/A |
| | | LSP_DEFAULT | | | | | 127.0.0.1 |
| L> | 3.3.3.3/32 | | 1 | 52480 | 3 | N/A | xe5 |
| | | LSP_DEFAULT | | | | | 20.20.20.2 |
| L> | 1.1.1.1/32 | | 2 | 52481 | Nolabel | N/A | N/A |
| | | LSP_DEFAULT | | | | | 127.0.0.1 |
| L> | 30.30.30.0/24 | | 4 | 52483 | Nolabel | N/A | N/A |
| | | LSP_DEFAULT | | | | | 127.0.0.1 |
| L> | 40.40.40.0/24 | | 5 | 52484 | Nolabel | N/A | N/A |
| | | LSP_DEFAULT | | | | | 127.0.0.1 |

Pl#

Pl#show mpls ftn-table

```

Primary FTN entry with FEC: 1.1.1.1/32, id: 2, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 2
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: xe5, out label: 24965
Nexthop addr: 20.20.20.2      cross connect ix: 1, op code: Push

```

```

Primary FTN entry with FEC: 3.3.3.3/32, id: 5, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 8

```

```

Owner: N/A, Persistent: NO, Admin Status: Up, Oper Status: Up
Out-segment with ix: 8, owner: N/A, Stale: NO, out intf: xe5, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 5, op code: Push

```

```

Primary FTN entry with FEC: 4.4.4.4/32, id: 3, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: xe5, out label: 24966
Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

```

```

Primary FTN entry with FEC: 30.30.30.0/24, id: 4, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 6, owner: LDP, Stale: NO, out intf: xe5, out label: 24967
Nexthop addr: 20.20.20.2          cross connect ix: 4, op code: Push

```

```

Primary FTN entry with FEC: 40.40.40.0/24, id: 1, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 8
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 8, owner: N/A, Stale: NO, out intf: xe5, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 5, op code: Push

```

```

Pl#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

```

```

IP Route Table for VRF "default"
O      1.1.1.1/32 [110/53] via 20.20.20.2, xe5, 00:03:44
C      2.2.2.2/32 is directly connected, lo, 00:14:13
O      3.3.3.3/32 [110/2] via 20.20.20.2, xe5, 00:12:51
O      4.4.4.4/32 [110/52] via 20.20.20.2, xe5, 00:03:44
C      20.20.20.0/24 is directly connected, xe5, 00:13:46
O      30.30.30.0/24 [110/52] via 20.20.20.2, xe5, 00:03:44
O      40.40.40.0/24 [110/51] via 20.20.20.2, xe5, 00:03:44
C      127.0.0.0/8 is directly connected, lo, 00:14:13

```

Bring up the link and check same labels reused.

| | |
|-------------------------|-----------------------|
| (config)#interface xe12 | Enter interface mode. |
| (config-if)#no shutdown | Shutdown the link. |
| (config)#commit | Commit. |

Up link Validation

```

PE1#show ldp session
Peer IP Address      IF Name    My Role    State    KeepAlive  UpTime

```

```

4.4.4.4          xe6      Passive  OPERATIONAL  30    00:14:55
2.2.2.2          xe12     Passive  OPERATIONAL  30    00:15:41
PE1#show ldp targeted-peers
IP Address      Interface
2.2.2.2         xe12
4.4.4.4         xe6

```

```

PE1#show ldp session 2.2.2.2
Session state      : OPERATIONAL
Session role       : Passive
TCP Connection     : Established
IP Address for TCP : 2.2.2.2
Interface being used : xe12
Peer LDP ID        : 2.2.2.2:0
Peer LDP Password   : Not Set
Adjacencies        : 10.10.10.2
                   2.2.2.2
Advertisement mode   : Downstream Unsolicited
Label retention mode : Liberal
Graceful Restart    : Not Capable
Keepalive Timeout   : 30
Reconnect Interval  : 15
Session protection   : Ready
Address List received : 2.2.2.2
                   10.10.10.2
                   20.20.20.1
                   254.128.0.0

```

```

Received Labels :   Fec          Label          Maps To
                  IPV4:3.3.3.3/32    52480         24966
                  IPV4:20.20.20.0/24 impl-null     24967
                  IPV4:10.10.10.0/24 impl-null     none
                  IPV4:2.2.2.2/32    impl-null     24965
Sent Labels :      Fec          Label          Maps To
                  IPV4:40.40.40.0/24 24961         impl-null
                  IPV4:4.4.4.4/32    24960         impl-null
                  IPV4:30.30.30.0/24 impl-null     none
                  IPV4:10.10.10.0/24 impl-null     none
                  IPV4:1.1.1.1/32    impl-null     none

```

```

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport

```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|--------|------------|-----------|-----|----------|-----------|------|
| Intf | ELC Nextthop | | | | | | | |
| L> | 2.2.2.2/32 | 3 | 9 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe12 | No | 10.10.10.2 | | | | | |
| L> | 3.3.3.3/32 | 1 | 2 | - | Yes | LSP_ | | |
| DEFAULT | 52480 xe12 | No | 10.10.10.2 | | | | | |
| L> | 4.4.4.4/32 | 4 | 7 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe6 | No | 30.30.30.2 | | | | | |
| L> | 20.20.20.0/24 | 2 | 4 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe12 | No | 10.10.10.2 | | | | | |
| L> | 40.40.40.0/24 | 5 | 8 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe6 | No | 30.30.30.2 | | | | | |

```

PE1#
PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|--------|----------|-----------|---------|------|
| Intf/VRF | Nextthop | | LSP-Type | | | |

```

L> 2.2.2.2/32 9 24965 3 N/A xe12 10.10.10.2
    LSP_DEFAULT
L> 40.40.40.0/24 2 24961 3 N/A xe6 30.30.30.2
    LSP_DEFAULT
L> 4.4.4.4/32 1 24960 3 N/A xe6 30.30.30.2
    LSP_DEFAULT
L> 3.3.3.3/32 10 24966 52480 N/A xe12 10.10.10.2
    LSP_DEFAULT
L> 20.20.20.0/24 11 24967 3 N/A xe12 10.10.10.2
    LSP_DEFAULT
PE1#show mpls ftn-table
Primary FTN entry with FEC: 2.2.2.2/32, id: 3, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 10
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 10, owner: N/A, Stale: NO, out intf: xe12, out label: 3
Nexthop addr: 10.10.10.2 cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 1, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 11
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 11, owner: LDP, Stale: NO, out intf: xe12, out label: 52480
Nexthop addr: 10.10.10.2 cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 4, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 6, owner: N/A, Stale: NO, out intf: xe6, out label: 3
Nexthop addr: 30.30.30.2 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 20.20.20.0/24, id: 2, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 10
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 10, owner: N/A, Stale: NO, out intf: xe12, out label: 3
Nexthop addr: 10.10.10.2 cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 40.40.40.0/24, id: 5, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 6
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 6, owner: N/A, Stale: NO, out intf: xe6, out label: 3
Nexthop addr: 30.30.30.2 cross connect ix: 4, op code: Push

P1#show ldp session
Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
3.3.3.3              xe5        Passive    OPERATIONAL 30      00:15:30
1.1.1.1              xe48       Active     OPERATIONAL 30      00:16:33

P1#show ldp targeted-peers
IP Address      Interface
1.1.1.1         xe48
3.3.3.3         xe5

P1#show ldp session 1.1.1.1
Session state      : OPERATIONAL

```

```

Session role           : Active
TCP Connection         : Established
IP Address for TCP     : 1.1.1.1
Interface being used   : xe48
Peer LDP ID           : 1.1.1.1:0
Peer LDP Password      : Not Set
Adjacencies           : 10.10.10.1
                       : 1.1.1.1
Advertisement mode      : Downstream Unsolicited
Label retention mode    : Liberal
Graceful Restart       : Not Capable
Keepalive Timeout      : 30
Reconnect Interval     : 15
Session protection     : Ready
Address List received  : 1.1.1.1
                       : 10.10.10.1
                       : 30.30.30.1
                       : 254.128.0.0

Received Labels :      Fec          Label          Maps To
                  IPV4:4.4.4.4/32    24960          52482
                  IPV4:40.40.40.0/24 24961          52484
                  IPV4:30.30.30.0/24 impl-null       52483
                  IPV4:10.10.10.0/24 impl-null       none
                  IPV4:1.1.1.1/32    impl-null       52481

Sent Labels :      Fec          Label          Maps To
                  IPV4:3.3.3.3/32    52480          impl-null
                  IPV4:20.20.20.0/24 impl-null       none
                  IPV4:10.10.10.0/24 impl-null       none
                  IPV4:2.2.2.2/32    impl-null       none

```

Pl#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|---------|------------|-----------|-----|----------|-----------|------|
| Intf | ELC | Nexthop | | | | | | |
| L> | 1.1.1.1/32 | 2 | 3 | - | Yes | LSP_ | | |
| DEFAULT | 3 | No | 10.10.10.1 | | | | | |
| L> | 3.3.3.3/32 | 5 | 9 | - | Yes | LSP_ | | |
| DEFAULT | 3 | No | 20.20.20.2 | | | | | |
| L> | 4.4.4.4/32 | 3 | 5 | - | Yes | LSP_ | | |
| DEFAULT | 24960 | No | 10.10.10.1 | | | | | |
| L> | 30.30.30.0/24 | 4 | 7 | - | Yes | LSP_ | | |
| DEFAULT | 3 | No | 10.10.10.1 | | | | | |
| L> | 40.40.40.0/24 | 1 | 1 | - | Yes | LSP_ | | |
| DEFAULT | 24961 | No | 10.10.10.1 | | | | | |

Pl#show mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM
 K - CLI ILM, T - MPLS-TP, s - Stitched ILM
 S - SNMP, L - LDP, R - RSVP, C - CRLDP
 B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
 U - unknown

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|--------|----------|-----------|---------|-----------------|
| Intf/VRF | Nexthop | | LSP-Type | | | |
| L> | 4.4.4.4/32 | 3 | 52482 | 24960 | N/A | xe48 10.10.10.1 |
| | LSP_DEFAULT | | | | | |
| L> | 1.1.1.1/32 | 2 | 52481 | 3 | N/A | xe48 10.10.10.1 |
| | LSP_DEFAULT | | | | | |
| L> | 3.3.3.3/32 | 1 | 52480 | 3 | N/A | xe5 20.20.20.2 |
| | LSP_DEFAULT | | | | | |
| L> | 40.40.40.0/24 | 5 | 52484 | 24961 | N/A | xe48 10.10.10.1 |
| | LSP_DEFAULT | | | | | |
| L> | 30.30.30.0/24 | 4 | 52483 | 3 | N/A | xe48 10.10.10.1 |
| | LSP_DEFAULT | | | | | |

```
Pl#show mpls ftn-table
Primary FTN entry with FEC: 1.1.1.1/32, id: 2, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 10
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 10, owner: N/A, Stale: NO, out intf: xe48, out label: 3
  Nexthop addr: 10.10.10.1      cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 5, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 8
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 8, owner: N/A, Stale: NO, out intf: xe5, out label: 3
  Nexthop addr: 20.20.20.2      cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 3, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 11
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 11, owner: LDP, Stale: NO, out intf: xe48, out label: 24960
  Nexthop addr: 10.10.10.1      cross connect ix: 7, op code: Push

Primary FTN entry with FEC: 30.30.30.0/24, id: 4, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 10
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 10, owner: N/A, Stale: NO, out intf: xe48, out label: 3
  Nexthop addr: 10.10.10.1      cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 40.40.40.0/24, id: 1, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 12
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 12, owner: LDP, Stale: NO, out intf: xe48, out label: 24961
  Nexthop addr: 10.10.10.1      cross connect ix: 8, op code: Push
```

LDP Inter-Area Configuration

This chapter contains LDP inter-area configuration examples.

Overview

Provider-based MPLS (Multiprotocol Label Switching) networks are expanding with the success of Layer 3 Virtual Private Networks and the new deployments of Layer 2 VPNs. Service providers MPLS backbones are significantly growing both in terms of density with the addition of Provider Edge (PE) routers to connect new customers and in terms of footprint as traditional Layer 2 aggregation networks may be replaced by IP/MPLS networks. As a consequence, many providers need to introduce IGP areas. Inter-area LSPs (that is, LSPs that traverse at least two IGP areas) are required to ensure MPLS connectivity between PEs located in distinct IGP areas.

On a large MPLS networks, multiple IGP areas need to be configured for flexible network deployment and fast route convergence. When advertising routes between IGP areas, to prevent a large number of routes from consuming too many resources, an Area Border Router (ABR) needs to aggregate the routes in the area and advertises the aggregated route to the neighboring IGP areas. By default, when establishing LSPs, LDP searches the routing table for the route that exactly matches the FEC in the received Label Mapping message.

The LDP inter-area feature provides a longest-match label mapping procedure where a label is used if the Forwarding Equivalence Class (FEC) matches an entry in the Routing Information Base (RIB). Matching is defined by an IP longest-match search and does not mandate an exact match.

Configure LDP Inter-Area

The LDP Inter-Area configuration process can be divided into the following tasks:

1. Enable label-switching on the interface in NSM.
2. Enable LDP on an interface in LDP.
3. Run an IGP (Internal Gateway Protocol) such as OSPF or ISIS to distribute reachability information within the MPLS cloud.
4. Give the `inter-area-lsp` command in LDP router mode to enable creation of inter-area LSPs.

To configure a summary route on an ABR, configure two networks such R4 and R5 in [Figure 7](#), so that two prefixes are summarized. For example, 4.4.4.0/24 is a summary route for the networks 4.4.4.1/32 and 4.4.4.2/32.


Give the `area-range` command on the routers to summarize the networks. For example: `area 2 range 4.4.4.0/24`.



Note: LDP Downstream-On-Demand not supported with inter-area configuration. LDP Inter-Area Graceful Restart (GR) is not supported in this release.


Topology

Figure 7. LDP inter-area topology




PE1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter loopback interface mode. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1 |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. |

| | |
|---|---|
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#targeted-peer ipv4 4.4.4.1 | Configure targeted peer. |
| (config-router-targeted-peer)#targeted-peer ipv4 4.4.4.2 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 10.1.1.1/24 | Set the IP address of the interface to 10.1.1.1/24 |
| (config)#router ospf 1 | Configure the routing process and specify the Process ID (1). |
| (config-router)#network 1.1.1.1/32 area 0 (config-router)#network 10.1.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter LDP router mode. |
| (config-router)#inter-area-lsp | Enable the inter-area-LSP command. |
| (config-router)#commit | Commit the transaction. |


P1

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 2.2.2.1 | Set the router ID to IP address 2.2.2.1 |
| (config-router)#transport-address ipv4 2.2.2.1 | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. <div>  Note: It is preferable to use the loopback address as the transport address. </div> |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe2 | Enter interface mode. |

| | |
|---|---|
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter loopback interface mode. |
| (config-if)#ip address 2.2.2.1/32 | Set the IP address of the loopback interface to 2.2.2.1/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 10.1.1.2/24 | Set the IP address of the interface to 10.1.1.2/24 |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#ip address 20.1.1.1/24 | Set the IP address of the interface to 20.1.1.1/24 |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the routing process and specify the Process ID (1). |
| (config-router)#network 2.2.2.1/32 area 0 (config-router)#network 10.1.1.0/24 area 0 (config-router)#network 20.1.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter LDP router mode. |
| (config-router)#inter-area-lsp | Enable the inter-area-LSP command. |
| (config-router)#commit | Commit the transaction. |


P2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter loopback interface mode. |
| (config-if)#ip address 3.3.3.1/32 secondary | Set the IP address of the loopback interface to 3.3.3.1/32. |
| (config-if)#exit | Exit interface mode. |

| | |
|---|--|
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ip address 20.1.1.2/24 | Set the IP address of the interface to 20.1.1.2/24 |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#ip address 30.1.1.1/24 | Set the IP address of the interface to 30.1.1.1/24 |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe3 | Specify the interface (xe3) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe1. |
| (config-if)#ip address 30.1.2.1/24 | Set the IP address of the interface to 30.1.2.1/24 |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the routing process and specify the Process ID (1). |
| (config-router)#network 3.3.3.1/32 area 0 (config-router)#network 20.1.1.0/24 area 0 (config-router)#network 30.1.1.0/24 area 2 (config-router)#network 30.1.2.0/24 area 2 | Define the interface on which OSPF runs and associate the area ID (0 and 2) respectively with the interface. |
| (config)#router ldp | Enter LDP router mode. |
| (config-router)#transport-address ipv4 3.3.3.1 | <p>Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run.</p> <div>  Note: Note: It is preferable to use the loopback address as the transport address. </div> |
| (config-router)#router-id 3.3.3.1 | Set the router ID to IP address 3.3.3.1 |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe3 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe3. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the routing process and specify the |


| | |
|---|---|
| | Process ID (1). |
| (config-router)#area 2 range 4.4.4.0/24 | Configure the summary route with range command. |
| (config-router)#commit | Commit the transaction. |

PE2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter loopback interface mode. |
| (config-if)#ip address 4.4.4.1/32 secondary | Set the IP address of the loopback interface to 4.4.4.1/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 30.1.1.2/24 | Set the IP address of the interface to 30.1.1.2/24 |
| (config)#router ospf 1 | Configure the routing process and specify the Process ID (1). |
| (config-router)#network 4.4.4.1/32 area 2 (config-router)#network 30.1.1.0/24 area 2 | Define the interface on which OSPF runs and associate the area ID (2) with the interface. |
| (config)#router ldp | Enter LDP router mode. |
| (config-router)#router-id 4.4.4.1 | Set the router ID to IP address to 4.4.4.1 |
| (config-router)#transport-address ipv4 4.4.4.1 | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. |
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| (config-if)#commit | Commit the transaction. |

PE3

| | |
|-----------------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#interface xe3 | Specify the interface(xe3) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe2. |

| | |
|---|---|
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter loopback interface mode. |
| (config-if)#ip address 4.4.4.2/32 secondary | Set the IP address of the loopback interface to 4.4.4.2/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 30.1.2.2/24 | Set the IP address of the interface to 30.1.2.2/24 |
| (config)#router ospf 1 | Configure the routing process and specify the Process ID (1). |
| (config-router)#network 4.4.4.2/32 area 2 (config-router)#network 30.1.2.0/24 area 2 | Define the interface on which OSPF runs and associate the area ID (2) with the interface. |
| (config)#router ldp | Enter LDP router mode. |
| (config-router)#router-id 4.4.4.2 | Set the router ID to IP address to 4.4.4.2 |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. |
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe3 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe3. |
| (config-if)#commit | Commit the transaction. |

Validation

PE1

```

PE1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
C      1.1.1.1/32 is directly connected, lo, installed 00:30:21, last update 00:30:21 ago
O      2.2.2.1/32 [110/2] via 10.1.1.2, xe5, installed 00:26:16, last update 00:26:16 ago
O      3.3.3.1/32 [110/3] via 10.1.1.2, xe5, installed 00:16:08, last update 00:16:08 ago
O IA   4.4.4.0/24 [110/4] via 10.1.1.2, xe5, installed 00:13:45, last update 00:13:45 ago
C      10.1.1.0/24 is directly connected, xe5, installed 00:30:21, last update 00:30:21 ago
O      20.1.1.0/24 [110/2] via 10.1.1.2, xe5, installed 00:23:26, last update 00:23:26 ago
O IA   30.1.1.0/24 [110/3] via 10.1.1.2, xe5, installed 00:16:08, last update 00:16:08 ago
O IA   30.1.2.0/24 [110/3] via 10.1.1.2, xe5, installed 00:16:08, last update 00:16:08 ago

```

C 127.0.0.0/8 is directly connected, lo, installed 20:19:48, last update 20:19:48 ago

Gateway of last resort is not set

PE1#show ip ospf neighbor

Total number of full neighbors: 1

OSPF process 1 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 2.2.2.1 | 1 | Full/Backup | 00:00:33 | 10.1.1.2 | xe5 | 0 |

PE1#show ip ospf interface

lo is up, line protocol is up

Internet Address 1.1.1.1/32, Area 0.0.0.0, MTU 16436

Process ID 1, VRF (default), Router ID 1.1.1.1, Network Type LOOPBACK, Cost: 1

Transmit Delay is 1 sec, State Loopback, TE Metric 1

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

xe5 is up, line protocol is up

Internet Address 10.1.1.1/24, Area 0.0.0.0, MTU 1500

Process ID 1, VRF (default), Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1

Transmit Delay is 1 sec, State DR, Priority 1, TE Metric 1

Designated Router (ID) 1.1.1.1, Interface Address 10.1.1.1

Backup Designated Router (ID) 2.2.2.1, Interface Address 10.1.1.2

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

Hello due in 00:00:00

Neighbor Count is 1, Adjacent neighbor count is 1

Suppress hello for 0 neighbor(s)

Hello received 168 sent 190, DD received 3 sent 4

LS-Req received 1 sent 1, LS-Upd received 15 sent 6

LS-Ack received 4 sent 12, Discarded 0

No authentication

PE1#show ldp session

Codes: m - MD5 password is not set/unset.

g - GR configuration not set/unset.

t - TCP MSS not set/unset.

Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 4.4.4.1 | xe5 | Passive | OPERATIONAL | 30 | 00:04:12 |
| | 4.4.4.2 | xe5 | Passive | OPERATIONAL | 30 | 00:03:27 |
| | 2.2.2.1 | xe5 | Passive | OPERATIONAL | 30 | 00:27:57 |

PE1#show ldp inter-area-fecs

LSR codes : E/N - LSR is egress/non-egress for this FEC,

L - LSR received a label for this FEC,

> - LSR will use this route for the FEC

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|----------------------------------|------|---------|-----------|-----|--------------|
| Matching RIB prefix - 4.4.4.0/24 | | | | | |
| 4.4.4.1/32 | NL> | 2.2.2.1 | 24325 | No | 10.1.1.2 |
| 4.4.4.2/32 | NL> | 2.2.2.1 | 24326 | No | 10.1.1.2 |

PE1#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup

B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,

L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,

U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|------------|---------|----------|-----------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 2.2.2.1/32 | 1 | 2 | - | - | - | - |

```

- N/A 00:28:44
1 - 10.1.1.2 - -
L> Yes 3 xe5 No 10.1.1.2 - -
3.3.3.1/32 3 4 - - -
- N/A 00:19:10
3 -
L> Yes 24321 xe5 No 10.1.1.2 - -
4.4.4.0/24 6 10 - - -
- N/A 00:16:47
9 -
L> Yes 24324 xe5 No 10.1.1.2 - -
4.4.4.1/32 7 12 - - -
- N/A 00:16:47
11 -
L> Yes 24325 xe5 No 10.1.1.2 - -
4.4.4.2/32 8 14 - - -
- N/A 00:14:29
13 -
L> Yes 24326 xe5 No 10.1.1.2 - -
20.1.1.0/24 2 2 - - -
- N/A 00:26:28
1 -
L> Yes 3 xe5 No 10.1.1.2 - -
30.1.1.0/24 4 6 - - -
- N/A 00:19:10
5 -
L> Yes 24322 xe5 No 10.1.1.2 - -
30.1.2.0/24 5 8 - - -
- N/A 00:19:10
7 -
Yes 24323 xe5 No 10.1.1.2 -

```

P1

```

Pl#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN,
v - vrf leaked
* - candidate default

IP Route Table for VRF "default"
O 1.1.1.1/32 [110/2] via 10.1.1.1, xe15, installed 00:31:33, last update 00:31:33 ago
C 2.2.2.1/32 is directly connected, lo, installed 00:31:58, last update 00:31:58 ago
O 3.3.3.1/32 [110/2] via 20.1.1.2, xe12, installed 00:21:29, last update 00:21:29 ago
O IA 4.4.4.0/24 [110/3] via 20.1.1.2, xe12, installed 00:19:06, last update 00:19:06 ago
C 10.1.1.0/24 is directly connected, xe15, installed 00:31:58, last update 00:31:58 ago
C 20.1.1.0/24 is directly connected, xe12, installed 00:28:47, last update 00:28:47 ago
O IA 30.1.1.0/24 [110/2] via 20.1.1.2, xe12, installed 00:21:29, last update 00:21:29 ago
O IA 30.1.2.0/24 [110/2] via 20.1.1.2, xe12, installed 00:21:29, last update 00:21:29 ago
C 127.0.0.0/8 is directly connected, lo, installed 20:25:32, last update 20:25:32 ago

Gateway of last resort is not set

Pl#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID Pri State Dead Time Address Interface Instance ID
1.1.1.1 1 Full/DR 00:00:29 10.1.1.1 xe15 0
3.3.3.1 1 Full/Backup 00:00:38 20.1.1.2 xe12 0

Pl#show ip ospf interface

```

```

lo is up, line protocol is up
Internet Address 2.2.2.1/32, Area 0.0.0.0, MTU 16436
Process ID 1, VRF (default), Router ID 2.2.2.1, Network Type LOOPBACK, Cost: 1
Transmit Delay is 1 sec, State Loopback, TE Metric 1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
xe12 is up, line protocol is up
Internet Address 20.1.1.1/24, Area 0.0.0.0, MTU 1500
Process ID 1, VRF (default), Router ID 2.2.2.1, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DR, Priority 1, TE Metric 1
Designated Router (ID) 2.2.2.1, Interface Address 20.1.1.1
Backup Designated Router (ID) 3.3.3.1, Interface Address 20.1.1.2
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:06
Neighbor Count is 1, Adjacent neighbor count is 1
Suppress hello for 0 neighbor(s)
Hello received 131 sent 179, DD received 3 sent 4
LS-Req received 1 sent 1, LS-Upd received 4 sent 7
LS-Ack received 5 sent 4, Discarded 0
No authentication
xe15 is up, line protocol is up
Internet Address 10.1.1.2/24, Area 0.0.0.0, MTU 1500
Process ID 1, VRF (default), Router ID 2.2.2.1, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State Backup, Priority 1, TE Metric 1
Designated Router (ID) 1.1.1.1, Interface Address 10.1.1.1
Backup Designated Router (ID) 2.2.2.1, Interface Address 10.1.1.2
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:03
Neighbor Count is 1, Adjacent neighbor count is 1
Suppress hello for 0 neighbor(s)
Hello received 190 sent 191, DD received 4 sent 3
LS-Req received 1 sent 1, LS-Upd received 7 sent 15
LS-Ack received 11 sent 5, Discarded 0
No authentication

```

Pl#show ldp session

```

Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 3.3.3.1 | xe12 | Passive | OPERATIONAL | 30 | 00:25:39 |
| | 1.1.1.1 | xe15 | Active | OPERATIONAL | 30 | 00:35:53 |

Pl#show ldp inter-area-fecs

```

LSR codes : E/N - LSR is egress/non-egress for this FEC,
           L - LSR received a label for this FEC,
           > - LSR will use this route for the FEC

```

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|----------------------------------|------|---------|-----------|-----|--------------|
| Matching RIB prefix - 4.4.4.0/24 | | | | | |
| 4.4.4.1/32 | NL | 1.1.1.1 | 24326 | No | no nexthop |
| | NL> | 3.3.3.1 | 24323 | No | 20.1.1.2 |
| 4.4.4.2/32 | NL | 1.1.1.1 | 24327 | No | no nexthop |
| | NL> | 3.3.3.1 | 24324 | No | 20.1.1.2 |

Pl#show mpls forwarding-table

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|-----|---------|----------|-----------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |


```

L> 1.1.1.1/32      1      2      -      -      -      -      -
    -             N/A      00:31:06
                        1
    Yes 3          xe15      No      10.1.1.1      -      -
L> 3.3.3.1/32      2      5      -      -      -      -      -
    -             N/A      00:20:50
                        4
    Yes 3          xe12      No      20.1.1.2      -      -
L> 4.4.4.0/24      6      5      -      -      -      -      -
    -             N/A      00:19:08
                        4
    Yes 3          xe12      No      20.1.1.2      -      -
L> 4.4.4.1/32      5      8      -      -      -      -      -
    -             N/A      00:19:08
                        7
    Yes 24323      xe12      No      20.1.1.2      -      -
L> 4.4.4.2/32      7      11     -      -      -      -      -
    -             N/A      00:16:50
                        10
    Yes 24324      xe12      No      20.1.1.2      -      -
L> 30.1.1.0/24     3      5      -      -      -      -      -
    -             N/A      00:20:50
                        4
    Yes 3          xe12      No      20.1.1.2      -      -
L> 30.1.2.0/24     4      5      -      -      -      -      -
    -             N/A      00:20:50
                        4
    Yes 3          xe12      No      20.1.1.2      -      -

```

P2

```
P2#show ip ospf neighbor
```

```
Total number of full neighbors: 3
```

```
OSPF process 1 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 2.2.2.1 | 1 | Full/DR | 00:00:31 | 20.1.1.1 | xe3 | 0 |
| 4.4.4.1 | 1 | Full/Backup | 00:00:33 | 30.1.1.2 | xe11 | 0 |
| 4.4.4.2 | 1 | Full/Backup | 00:00:38 | 30.1.2.2 | xe6 | 0 |

```
P2#show ip route
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
```

```
O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
```

```
ia - IS-IS inter area, E - EVPN,
```

```
v - vrf leaked
```

```
* - candidate default
```

```
IP Route Table for VRF "default"
```

```

O      1.1.1.1/32 [110/3] via 20.1.1.1, xe3, installed 00:23:49, last update 00:23:49 ago
O      2.2.2.1/32 [110/2] via 20.1.1.1, xe3, installed 00:23:49, last update 00:23:49 ago
C      3.3.3.1/32 is directly connected, lo, installed 00:24:11, last update 00:24:11 ago
O      4.4.4.0/24 [110/0] is a summary, Null, installed 00:21:32, last update 00:21:32 ago
O      4.4.4.1/32 [110/2] via 30.1.1.2, xe11, installed 00:21:32, last update 00:21:32 ago
O      4.4.4.2/32 [110/2] via 30.1.2.2, xe6, installed 00:19:16, last update 00:19:16 ago
O      10.1.1.0/24 [110/2] via 20.1.1.1, xe3, installed 00:23:49, last update 00:23:49 ago
C      20.1.1.0/24 is directly connected, xe3, installed 00:24:11, last update 00:24:11 ago
C      30.1.1.0/24 is directly connected, xe11, installed 00:24:11, last update 00:24:11 ago
C      30.1.2.0/24 is directly connected, xe6, installed 00:24:11, last update 00:24:11 ago
C      127.0.0.0/8 is directly connected, lo, installed 00:31:38, last update 00:31:38 ago

```

```
P2#show ip ospf interface
```

```
lo is up, line protocol is up
```

```

Internet Address 3.3.3.1/32, Area 0.0.0.0, MTU 16436
Process ID 1, VRF (default), Router ID 3.3.3.1, Network Type LOOPBACK, Cost: 1
Transmit Delay is 1 sec, State Loopback, TE Metric 1
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
xe3 is up, line protocol is up
Internet Address 20.1.1.2/24, Area 0.0.0.0, MTU 1500
Process ID 1, VRF (default), Router ID 3.3.3.1, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State Backup, Priority 1, TE Metric 1
Designated Router (ID) 2.2.2.1, Interface Address 20.1.1.1
Backup Designated Router (ID) 3.3.3.1, Interface Address 20.1.1.2
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:06
Neighbor Count is 1, Adjacent neighbor count is 1
Suppress hello for 0 neighbor(s)
Hello received 147 sent 148, DD received 4 sent 3
LS-Req received 1 sent 1, LS-Upd received 8 sent 4
LS-Ack received 3 sent 6, Discarded 0
No authentication
xe6 is up, line protocol is up
Internet Address 30.1.2.1/24, Area 0.0.0.2, MTU 1500
Process ID 1, VRF (default), Router ID 3.3.3.1, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DR, Priority 1, TE Metric 1
Designated Router (ID) 3.3.3.1, Interface Address 30.1.2.1
Backup Designated Router (ID) 4.4.4.2, Interface Address 30.1.2.2
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:01
Neighbor Count is 1, Adjacent neighbor count is 1
Suppress hello for 0 neighbor(s)
Hello received 120 sent 147, DD received 3 sent 4
LS-Req received 1 sent 1, LS-Upd received 3 sent 7
LS-Ack received 6 sent 3, Discarded 0
No authentication
xe11 is up, line protocol is up
Internet Address 30.1.1.1/24, Area 0.0.0.2, MTU 1500
Process ID 1, VRF (default), Router ID 3.3.3.1, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DR, Priority 1, TE Metric 1
Designated Router (ID) 3.3.3.1, Interface Address 30.1.1.1
Backup Designated Router (ID) 4.4.4.1, Interface Address 30.1.1.2
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:07
Neighbor Count is 1, Adjacent neighbor count is 1
Suppress hello for 0 neighbor(s)
Hello received 134 sent 147, DD received 3 sent 4
LS-Req received 1 sent 1, LS-Upd received 4 sent 15
LS-Ack received 11 sent 3, Discarded 0
No authentication

```

P2#show ldp session

```

Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 4.4.4.2 | xe6 | Passive | OPERATIONAL | 30 | 00:10:06 |
| | 2.2.2.1 | xe3 | Active | OPERATIONAL | 30 | 00:24:39 |
| | 4.4.4.1 | xe11 | Passive | OPERATIONAL | 30 | 00:22:21 |

P2#show mpls forwarding-table

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|-------------|---------|----------|-----------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 1 | 2 | - | - | - | - |
| | - | N/A | 00:23:23 | | | | |
| | Yes | 24320 | xe3 | No | | 20.1.1.1 | - |
| L> | 2.2.2.1/32 | 2 | 4 | - | - | - | - |
| | - | N/A | 00:23:23 | | | | |
| | Yes | 3 | xe3 | No | | 20.1.1.1 | - |
| L> | 4.4.4.1/32 | 4 | 8 | - | - | - | - |
| | - | N/A | 00:21:07 | | | | |
| | Yes | 3 | xe11 | No | | 30.1.1.2 | - |
| L> | 4.4.4.2/32 | 5 | 11 | - | - | - | - |
| | - | N/A | 00:08:52 | | | | |
| | Yes | 3 | xe6 | No | | 30.1.2.2 | - |
| L> | 10.1.1.0/24 | 3 | 4 | - | - | - | - |
| | - | N/A | 00:23:23 | | | | |
| | Yes | 3 | xe3 | No | | 20.1.1.1 | - |

PE2

PE2#show ip route

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP

O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,

ia - IS-IS inter area, E - EVPN,

v - vrf leaked

* - candidate default

IP Route Table for VRF "default"

```

O IA      1.1.1.1/32 [110/4] via 30.1.1.1, xe11, installed 00:24:48, last update 00:24:48 ago
O IA      2.2.2.1/32 [110/3] via 30.1.1.1, xe11, installed 00:24:48, last update 00:24:48 ago
O IA      3.3.3.1/32 [110/2] via 30.1.1.1, xe11, installed 00:24:48, last update 00:24:48 ago
C         4.4.4.1/32 is directly connected, lo, installed 00:25:09, last update 00:25:09 ago
O         4.4.4.2/32 [110/3] via 30.1.1.1, xe11, installed 00:22:30, last update 00:22:30 ago
O IA      10.1.1.0/24 [110/3] via 30.1.1.1, xe11, installed 00:24:48, last update 00:24:48 ago
O IA      20.1.1.0/24 [110/2] via 30.1.1.1, xe11, installed 00:24:48, last update 00:24:48 ago
C         30.1.1.0/24 is directly connected, xe11, installed 00:25:09, last update 00:25:09 ago
O         30.1.2.0/24 [110/2] via 30.1.1.1, xe11, installed 00:24:48, last update 00:24:48 ago
C         127.0.0.0/8 is directly connected, lo, installed 00:34:25, last update 00:34:25 ago

```

Gateway of last resort is not set

PE2#show ip ospf neighbor

Total number of full neighbors: 1

OSPF process 1 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|----------|-----------|-------------|
| 3.3.3.1 | 1 | Full/DR | 00:00:36 | 30.1.1.1 | xe11 | 0 |

PE2#show ip ospf interface

lo is up, line protocol is up

Internet Address 4.4.4.1/32, Area 0.0.0.2, MTU 16436

Process ID 1, VRF (default), Router ID 4.4.4.1, Network Type LOOPBACK, Cost: 1

Transmit Delay is 1 sec, State Loopback, TE Metric 1

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

xe11 is up, line protocol is up

Internet Address 30.1.1.2/24, Area 0.0.0.2, MTU 1500

Process ID 1, VRF (default), Router ID 4.4.4.1, Network Type BROADCAST, Cost: 1

```

Transmit Delay is 1 sec, State Backup, Priority 1, TE Metric 1
Designated Router (ID) 3.3.3.1, Interface Address 30.1.1.1
Backup Designated Router (ID) 4.4.4.1, Interface Address 30.1.1.2
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:08
Neighbor Count is 1, Adjacent neighbor count is 1
Suppress hello for 0 neighbor(s)
Hello received 1052 sent 1053, DD received 4 sent 3
LS-Req received 1 sent 1, LS-Upd received 83 sent 20
LS-Ack received 19 sent 78, Discarded 0
No authentication

```

PE2#show ldp session

```

Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 3.3.3.1 | xe11 | Active | OPERATIONAL | 30 | 00:24:40 |
| | 1.1.1.1 | xe11 | Active | OPERATIONAL | 30 | 00:13:28 |
| | 4.4.4.2 | xe11 | Passive | OPERATIONAL | 30 | 00:12:38 |

PE2#show mpls forwarding-table

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|-------------|---------|----------|-----------|----------|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 1 | 2 | - | - | - | - |
| | - | N/A | 00:24:49 | | | | |
| | Yes | 24320 | xe11 | No | 30.1.1.1 | - | - |
| L> | 2.2.2.1/32 | 2 | 4 | - | - | - | - |
| | - | N/A | 00:24:49 | | | | |
| | Yes | 24321 | xe11 | No | 30.1.1.1 | - | - |
| L> | 3.3.3.1/32 | 3 | 6 | - | - | - | - |
| | - | N/A | 00:24:49 | | | | |
| | Yes | 3 | xe11 | No | 30.1.1.1 | - | - |
| L> | 4.4.4.2/32 | 7 | 10 | - | - | - | - |
| | - | N/A | 00:23:03 | | | | |
| | Yes | 24324 | xe11 | No | 30.1.1.1 | - | - |
| L> | 10.1.1.0/24 | 4 | 8 | - | - | - | - |
| | - | N/A | 00:24:49 | | | | |
| | Yes | 24322 | xe11 | No | 30.1.1.1 | - | - |
| L> | 20.1.1.0/24 | 5 | 6 | - | - | - | - |
| | - | N/A | 00:24:49 | | | | |
| | Yes | 3 | xe11 | No | 30.1.1.1 | - | - |
| L> | 30.1.2.0/24 | 6 | 6 | - | - | - | - |
| | - | N/A | 00:24:49 | | | | |
| | Yes | 3 | xe11 | No | 30.1.1.1 | - | - |

PE3

PE3#show ip route

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP

O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,

ia - IS-IS inter area, E - EVPN,

v - vrf leaked

* - candidate default

IP Route Table for VRF "default"

```
O IA      1.1.1.1/32 [110/4] via 30.1.2.1, xe1, installed 00:34:39, last update 00:34:39 ago
O IA      2.2.2.1/32 [110/3] via 30.1.2.1, xe1, installed 00:34:39, last update 00:34:39 ago
O IA      3.3.3.1/32 [110/2] via 30.1.2.1, xe1, installed 00:34:39, last update 00:34:39 ago
O         4.4.4.1/32 [110/3] via 30.1.2.1, xe1, installed 00:34:39, last update 00:34:39 ago
C         4.4.4.2/32 is directly connected, lo, installed 00:35:05, last update 00:35:05 ago
O IA      10.1.1.0/24 [110/3] via 30.1.2.1, xe1, installed 00:34:39, last update 00:34:39 ago
O IA      20.1.1.0/24 [110/2] via 30.1.2.1, xe1, installed 00:34:39, last update 00:34:39 ago
O         30.1.1.0/24 [110/2] via 30.1.2.1, xe1, installed 00:34:39, last update 00:34:39 ago
C         30.1.2.0/24 is directly connected, xe1, installed 00:35:05, last update 00:35:05 ago
C         127.0.0.0/8 is directly connected, lo, installed 20:42:59, last update 20:42:59 ago
```

Gateway of last resort is not set

PE3#show ip ospf neighbor

Total number of full neighbors: 1

OSPF process 1 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|----------|-----------|-------------|
| 3.3.3.1 | 1 | Full/DR | 00:00:35 | 30.1.2.1 | xe1 | 0 |

PE3#show ip ospf interface

lo is up, line protocol is up

Internet Address 4.4.4.2/32, Area 0.0.0.2, MTU 16436

Process ID 1, VRF (default), Router ID 4.4.4.2, Network Type LOOPBACK, Cost: 1

Transmit Delay is 1 sec, State Loopback, TE Metric 1

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

xe1 is up, line protocol is up

Internet Address 30.1.2.2/24, Area 0.0.0.2, MTU 1500

Process ID 1, VRF (default), Router ID 4.4.4.2, Network Type BROADCAST, Cost: 1

Transmit Delay is 1 sec, State Backup, Priority 1, TE Metric 1

Designated Router (ID) 3.3.3.1, Interface Address 30.1.2.1

Backup Designated Router (ID) 4.4.4.2, Interface Address 30.1.2.2

Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5

Hello due in 00:00:05

Neighbor Count is 1, Adjacent neighbor count is 1

Suppress hello for 0 neighbor(s)

Hello received 211 sent 212, DD received 4 sent 3

LS-Req received 1 sent 1, LS-Upd received 17 sent 6

LS-Ack received 5 sent 16, Discarded 0

No authentication

PE3#show ldp session

Codes: m - MD5 password is not set/unset.

g - GR configuration not set/unset.

t - TCP MSS not set/unset.

Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 3.3.3.1 | xe1 | Active | OPERATIONAL | 30 | 00:10:25 |
| | 4.4.4.1 | xe1 | Active | OPERATIONAL | 30 | 00:24:55 |
| | 1.1.1.1 | xe1 | Active | OPERATIONAL | 30 | 00:24:58 |

PE3#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup

B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,

L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhife-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|-------------|---------|----------|-----------|----------|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 1 | 21 | - | - | - | - |
| | - | N/A | 00:10:27 | - | - | - | - |
| | Yes | 24320 | xe1 | No | 30.1.2.1 | - | - |
| L> | 2.2.2.1/32 | 2 | 23 | - | - | - | - |
| | - | N/A | 00:10:27 | - | - | - | - |
| | Yes | 24321 | xe1 | No | 30.1.2.1 | - | - |
| L> | 3.3.3.1/32 | 3 | 24 | - | - | - | - |
| | - | N/A | 00:10:27 | - | - | - | - |
| | Yes | 3 | xe1 | No | 30.1.2.1 | - | - |
| L> | 4.4.4.1/32 | 4 | 26 | - | - | - | - |
| | - | N/A | 00:10:27 | - | - | - | - |
| | Yes | 24322 | xe1 | No | 30.1.2.1 | - | - |
| L> | 10.1.1.0/24 | 5 | 28 | - | - | - | - |
| | - | N/A | 00:10:27 | - | - | - | - |
| | Yes | 24324 | xe1 | No | 30.1.2.1 | - | - |
| L> | 20.1.1.0/24 | 6 | 24 | - | - | - | - |
| | - | N/A | 00:10:27 | - | - | - | - |
| | Yes | 3 | xe1 | No | 30.1.2.1 | - | - |
| L> | 30.1.1.0/24 | 7 | 24 | - | - | - | - |
| | - | N/A | 00:10:27 | - | - | - | - |
| | Yes | 3 | xe1 | No | 30.1.2.1 | - | - |

MPLS LDP-IGP Synchronization

This chapter provides configuration information to enable MPLS LDP-IGP Synchronization capabilities.

Overview

Multi-Protocol Label Switching (MPLS) Label Distribution Protocol (LDP) Interior Gateway Protocol (IGP) Synchronization ensures that LDP is fully established before the IGP path is used for switching. In certain networks, there is dependency on the edge-to-edge Label Switched Paths (LSPs) setup by the Label Distribution Protocol (LDP), e.g., networks that are used for Multi-Protocol Label Switching (MPLS) Virtual Private Network (VPN) applications. For such applications, it is not possible to rely on Internet Protocol (IP) forwarding if the MPLS LSP is not operating appropriately. Labeled traffic can be dropped due to presence of black holes in situations where the Interior Gateway Protocol (IGP) is operational on a link but LDP sessions are not up as the label distribution is not completed. While the link could still be used for IP forwarding, it is not useful for MPLS forwarding, for example, MPLS VPN applications or Border Gateway Protocol (BGP) route-free cores.

The MPLS LDP-IGP Synchronization feature ensures that the Label Distribution Protocol (LDP) is fully established before the Interior Gateway Protocol (IGP) path is used for packet forwarding. It is useful for cases in which the router is the ingress and the decision of whether to take the MPLS LSP or IGP path is decided there.

LDP-IGP synchronization is an interface level feature. It can be selectively enabled in the required interfaces. For each interface there are two commands available for synchronization, one each for IS-IS. Once configured the IGP saves the required information, and also notifies LDP. In between the IGP increases the link cost to maximum and sends advertisements to its peer. This discourages its peers from taking routes that pass via it.

When all LDP sessions hosted on the interface become operational, it sends a notification to the IGP. This is termed as LDP convergence. The IGP then advertises normal cost, so that all traffic now coming to the interface takes the MPLS LSP path established by LDP and not be IP routed.


Prerequisites

Only interfaces that are running Open Shortest Path First (OSPF) or Intermediate System-to-Intermediate System (IS-IS) processes are capable of LDP-IGP synchronization. The router must also be running LDP.

Topology

The sample topology diagram is applicable to all configurations in this chapter.

Figure 8. Sample Topology for LDP-IGP Synchronization



LDP-IGP Synchronization with OSPF

When IGP synchronization is enabled on OSPF-enabled interfaces, OSPF sends Maximum/Normal cost based on LDP session Down or Up state messages to interfaces until the hold-down-timer expires or synchronization is achieved.

Before configuring LDP-IGP synchronization, the NSM, OSPF and LDP configurations must be completed. The tables below contain examples of how this is done.

PE1 - NSM

| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#ip address 10.10.10.1/24 | Configure IPv4 address for xe11. |
| (config-if)#label-switching | Enable label switching on interface xe11. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| (config-if)#commit | Commit the transaction. |

PE1 - OSPF

| | |
|---|---|
| (config)#router ospf 100 | Configure the routing process and specify the Process ID 100. The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 1.1.1.1 | Configure OSPF router ID same as loopback IP address. |
| (config-router)#network 10.10.10.0/24 area 0 (config-router)#network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

PE1 - LDP

| | |
|--|--|
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1. |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |



Note: It is preferable to use the loopback address as the transport address.

| | |
|-----------------------------|------------------------------|
| (config-router)#exit | Exit router mode. |
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe11. |
| (config-if)#commit | Commit the transaction. |


PE2 - NSM

| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#ip address 10.10.10.2/24 | Configure IPv4 address for xe11. |
| (config-if)#label-switching | Enable label switching on interface xe11. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| (config-if)#commit | Commit the transaction. |

PE2 - OSPF

| | |
|---|---|
| (config)#router ospf 100 | Configure the routing process and specify the Process ID 100. The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 2.2.2.2 | Configure OSPF router ID same as loopback IP address. |
| (config-router)#network 10.10.10.0/24 area 0 (config-router)#network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

PE2 - LDP

| | |
|--|---|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2. |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| |  Note: It is preferable to use the loopback address as transport address. |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |

| | |
|-----------------------------|------------------------------|
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe11. |
| (config-if)#commit | Commit the transaction. |

Validation

PE1

```
PE1#show ip ospf neighbor
```

```
Total number of full neighbors: 1
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 2.2.2.2 | 1 | Full/Backup | 00:00:30 | 10.10.10.2 | xe11 | 0 |

```
PE1#show ip ospf interface brief
```

| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
|-----------|-----|---------|---------|------|----------|-----------|--------|
| lo | 100 | 0.0.0.0 | 1 | 1 | Loopback | 0 | Up |
| xe11 | 100 | 0.0.0.0 | 10032 | 1 | DR | 1 | Up |

```
PE1#show ip ospf interface
```

```
lo is up, line protocol is up
```

```
Internet Address 1.1.1.1/32, Area 0.0.0.0, MTU 16436
```

```
Process ID 100, VRF (default), Router ID 1.1.1.1, Network Type LOOPBACK, Cost: 1
```

```
Transmit Delay is 1 sec, State Loopback, TE Metric 1
```

```
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

```
xe11 is up, line protocol is up
```

```
Internet Address 10.10.10.1/24, Area 0.0.0.0, MTU 1500
```

```
Process ID 100, VRF (default), Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
```

```
Transmit Delay is 1 sec, State DR, Priority 1, TE Metric 1
```

```
Designated Router (ID) 1.1.1.1, Interface Address 10.10.10.1
```

```
Backup Designated Router (ID) 2.2.2.2, Interface Address 10.10.10.2
```

```
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
```

```
Hello due in 00:00:04
```

```
Neighbor Count is 1, Adjacent neighbor count is 1
```

```
Suppress hello for 0 neighbor(s)
```

```
Hello received 288 sent 304, DD received 3 sent 4
```

```
LS-Req received 1 sent 1, LS-Upd received 6 sent 9
```

```
LS-Ack received 7 sent 5, Discarded 0
```

```
No authentication
```

```
PE1#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
```

```
g - GR configuration not set/unset.
```

```
t - TCP MSS not set/unset.
```

```
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | xe11 | Passive | OPERATIONAL | 30 | 00:39:43 |

PE2

```
PE2#show ip ospf neighbor
```

```
Total number of full neighbors: 1
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|------------|-----------|-------------|
| 1.1.1.1 | 1 | Full/DR | 00:00:37 | 10.10.10.1 | xe11 | 0 |

```

PE2#show ip ospf interface brief
Interface      PID   Area      Intf ID   Cost   State      Neighbors  Status
lo             100   0.0.0.0    1         1     Loopback    0          Up
xe11          100   0.0.0.0   10011     1     Backup      1          Up

PE2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
1.1.1.1      xe11        Active     OPERATIONAL  30      00:41:38

PE2#show ip ospf interface
lo is up, line protocol is up
  Internet Address 2.2.2.2/32, Area 0.0.0.0, MTU 16436
  Process ID 100, VRF (default), Router ID 2.2.2.2, Network Type LOOPBACK, Cost: 1
  Transmit Delay is 1 sec, State Loopback, TE Metric 1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
xe11 is up, line protocol is up
  Internet Address 10.10.10.2/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 2.2.2.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State Backup, Priority 1, TE Metric 1
  Designated Router (ID) 1.1.1.1, Interface Address 10.10.10.1
  Backup Designated Router (ID) 2.2.2.2, Interface Address 10.10.10.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:04
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 291 sent 292, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 9 sent 6
  LS-Ack received 4 sent 7, Discarded 0
  No authentication


```

LDP-IGP Synchronization

Now that NSM, OSPF and LDP are all enabled, the LDP-IGP synchronization can be configured.

PE1

| | |
|---|---|
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#mpls ldp-igp sync ospf holddown-timer 500 | <p>Enable LDP-IGP Synchronization for xe11 belonging to an OSPF process and 500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> |

| | |
|---|--|
| |  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer. |
| <code>(config-if)#mpls ldp-igp sync-delay 60</code> | Configure time delay in seconds for notification of LDP convergence to IGP. This is not applicable for notification of non-convergence. Range is 5 to 60 seconds. This command is optional. LDP: This command is part of LDP Process. Default: If not configured the delay is 0 seconds. |
| <code>(config-if)#commit</code> | Commit the transaction. |

PE2

| | |
|--|---|
| <code>(config)#interface xe11</code> | Enter interface mode. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | Enable LDP-IGP Synchronization for interfaces (xe11) belonging to an OSPF process and 500 secs is Holddown-timer value for IGP to wait until LDP Converge. OSPF: This command is part of the OSPF Process. Note: Holddown-timer range is <1-2147483> seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer. |
| <code>(config-if)#mpls ldp-igp sync-delay 60</code> | Configure the time delay in seconds for the notification of LDP convergence to IGP. (This is not applicable for notification of non-convergence.) Range is 5 to 60 seconds. This command is optional. LDP: This command is part of LDP Process. Default: If not configured the delay is 0 seconds. |
| <code>(config-if)#commit</code> | Commit the transaction. |

PE1 Validation

When LDP IGP SYNC is Configured with hold-down and sync-delay timer

```

PE1#show ip ospf interface
lo is up, line protocol is up
  Internet Address 1.1.1.1/32, Area 0.0.0.0, MTU 16436
  Process ID 100, VRF (default), Router ID 1.1.1.1, Network Type LOOPBACK, Cost: 1
  Transmit Delay is 1 sec, State Loopback, TE Metric 1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
xe11 is up, line protocol is up
  Internet Address 10.10.10.1/24, Area 0.0.0.0, MTU 1500
  
```

```

Process ID 100, VRF (default), Router ID 1.1.1.1, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DR, Priority 1, TE Metric 1
LDP-OSPF Sync configured
  Holddown timer : 500 seconds, Remaining time = 0 seconds
Designated Router (ID) 1.1.1.1, Interface Address 10.10.10.1
Backup Designated Router (ID) 2.2.2.2, Interface Address 10.10.10.2
Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:04
Neighbor Count is 1, Adjacent neighbor count is 1
Suppress hello for 0 neighbor(s)
Hello received 262 sent 278, DD received 3 sent 4
LS-Req received 1 sent 1, LS-Upd received 6 sent 9
LS-Ack received 7 sent 5, Discarded 0
No authentication

```

```

PE1#show mpls ldp igp sync
Interface      Link-State LDP-Enabled IGP-Sync   Peer-IP      Session-State Sync-Delay-
Time/Remaining-Delay-Time
xe11           Up         Yes         Enabled    2.2.2.2      Achieved      60 sec / Not-Running

```

PE2 Validation

```

PE2#show ip ospf interface
lo is up, line protocol is up
  Internet Address 2.2.2.2/32, Area 0.0.0.0, MTU 16436
  Process ID 100, VRF (default), Router ID 2.2.2.2, Network Type LOOPBACK, Cost: 1
  Transmit Delay is 1 sec, State Loopback, TE Metric 1
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
xe11 is up, line protocol is up
  Internet Address 10.10.10.2/24, Area 0.0.0.0, MTU 1500
  Process ID 100, VRF (default), Router ID 2.2.2.2, Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State Backup, Priority 1, TE Metric 1
  LDP-OSPF Sync configured
    Holddown timer : 500 seconds, Remaining time = 0 seconds
  Designated Router (ID) 1.1.1.1, Interface Address 10.10.10.1
  Backup Designated Router (ID) 2.2.2.2, Interface Address 10.10.10.2
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:02
  Neighbor Count is 1, Adjacent neighbor count is 1
  Suppress hello for 0 neighbor(s)
  Hello received 266 sent 267, DD received 4 sent 3
  LS-Req received 1 sent 1, LS-Upd received 9 sent 6
  LS-Ack received 4 sent 7, Discarded 0
  No authentication

```

```

PE2#show mpls ldp igp sync
Interface      Link-State LDP-Enabled IGP-Sync   Peer-IP      Session-State Sync-Delay-
Time/Remaining-Delay-Time
xe11           Up         Yes         Enabled    1.1.1.1      Achieved      60 sec / Not-Running

```

LDP-IGP Synchronization with IS-IS

When IGP synchronization is enabled on an IS-IS enabled interfaces, IS-IS sends Maximum/Normal cost based on LDP session or Up state on interfaces until hold-down-timer expires or synchronization is achieved.

Before configuring LDP-IGP synchronization, the NSM, IS-IS and LDP configurations must be completed. The tables below contain examples of how this is done.

PE1 - NSM


| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#ip address 10.10.10.1/24 | Set the IP address of the xe11 to 10.10.10.1/24. |
| (config-if)#label-switching | Enable label switching on xe11. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| (config-if)#commit | Commit the transaction. |

PE1 - IS-IS

| | |
|---|--|
| (config)#router isis 1 | Configure the IS-IS routing instance and specify the TAG (1). The TAG should be a WORD - ISO routing area tag. |
| (config-router)#is-type level-1 | Define the IS to the specified level of routing for router. |
| (config-router)#net 49.0001.0000.0000.0001.00 | Configure the Network Entity Title (NET) for the instance. |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#isis circuit-type level-1 | Define the circuit type for the interface on which IS-IS runs and associate the level 1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode for the loopback interface (lo). |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#isis circuit-type level-1 | Define the circuit type for the interface on which IS-IS runs and associate the level 1. |
| (config-if)#commit | Commit the transaction. |

PE1 - LDP

| | |
|--|---|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1. |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address for IPV4 (for IPV6 use an IPV6 address) to use for a TCP session over |

| | |
|--|---|
| | <p>which LDP will run.</p> <div>  Note: It is preferable to use the loopback address as transport address. </div> |
| <code>(config-router)#exit</code> | Exit the Router mode and return to the Configure mode. |
| <code>(config)#interface xe11</code> | Enter interface mode. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe11. |
| <code>(config-if)#commit</code> | Commit the transaction. |

PE2 - NSM


| | |
|--|---|
| <code>#configure terminal</code> | Enter configuration mode |
| <code>(config)#interface xe11</code> | Enter interface mode. |
| <code>(config-if)#ip address 10.10.10.2/24</code> | Set the IP address of xe11 to 10.10.10.2/24 |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe11. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface lo</code> | Specify the loopback (lo) interface to be configured. |
| <code>(config-if)#ip address 2.2.2.2/32 secondary</code> | Set the IP address of the loopback interface to 2.2.2.2/32. |
| <code>(config-if)#commit</code> | Commit the transaction. |

PE2 - IS-IS

| | |
|--|---|
| <code>(config)#router isis 1</code> | Configure the IS-IS routing instance and specify the TAG as 1. The TAG should be a WORD - ISO routing area tag. |
| <code>(config-router)#is-type level-1</code> | Define the IS to the specified level of routing for router. |
| <code>(config-router)#net 49.0001.0000.0000.0002.00</code> | Configure the Network Entity Title (NET) for the instance. |
| <code>(config-router)#exit</code> | Exit the Router mode and return to the Configure mode. |
| <code>(config)#interface xe11</code> | Enter interface mode. |
| <code>(config-if)#ip router isis 1</code> | Configure IS-IS IPv4 routing on the interface with is-is tag instance 1. |
| <code>(config-if)#isis circuit-type level-1</code> | Define the circuit type for the interface on which IS-IS runs and associate the level type (1). |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface lo</code> | Enter interface mode for the loopback (lo) interface. |
| <code>(config-if)#ip router isis 1</code> | Configure IS-IS IPv4 routing on the interface with IS- |

| | |
|---------------------------------------|--|
| | IS tag instance 1. |
| (config-if)#isis circuit-type level-1 | Define the circuit type for the interface on which IS-IS runs and associate the level 1. |
| (config-if)#commit | Commit the transaction. |

PE2 - LDP

| | |
|--|---|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2. |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address for IPv4 (for IPv6 use an IPv6 address) to use for a TCP session over which LDP will run. <div>  Note: It is preferable to use the loopback address as transport address. </div> |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe11. |
| (config-if)#commit | Commit the transaction. |

Validation

```
PE1#show clns neighbors
```

```
Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1: VRF : default
System Id      Interface  SNPA          State Holdtime  Type Protocol
0000.0000.0002 xe11      04f8.f8a1.60f9 Up      22       L1   IS-IS
```

```
PE2#show clns neighbors
```

```
Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1: VRF : default
System Id      Interface  SNPA          State Holdtime  Type Protocol
0000.0000.0001 xe11      e8c5.7ad5.3d7a Up       7       L1   IS-IS
```

```
PE1#show clns is-neighbors
```

```
Tag 1: VRF : default
System Id      Interface  State  Type Priority  Circuit Id
0000.0000.0002 xe11      Up     L1    64        0000.0000.0001.01
```

```
PE2#show clns is-neighbors
```



```

Tag 1: VRF : default
System Id      Interface      State  Type Priority  Circuit Id
0000.0000.0001 xe11         Up     L1   64         0000.0000.0001.01

```

```

PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | xe11 | Passive | OPERATIONAL | 30 | 01:03:14 |

```

PE1#show ldp adjacency
Remote-Address  Local-Address  Mode          Intf-Name  Holdtime  LDP-Identifier
10.10.10.2      10.10.10.1    Interface     xe11       15        2.2.2.2:0

```

```

PE2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | xe11 | Active | OPERATIONAL | 30 | 01:03:50 |

```

PE2#show ldp adjacency
Remote-Address  Local-Address  Mode          Intf-Name  Holdtime  LDP-Identifier
10.10.10.1      10.10.10.2    Interface     xe11       15        1.1.1.1:0

```

```

PE1#show isis interface xe11
xe11 is up, line protocol is up
Routing Protocol: IS-IS (1)
  Network Type: Broadcast
  Circuit Type: level-1
  Local circuit ID: 0x01
  Extended Local circuit ID: 0x00002730
  Local SNPA: e8c5.7ad5.3d7a
  IP interface address:
    10.10.10.1/24
  IPv6 interface address:
    fe80::eac5:7aff:fed5:3d7a/64
  Level-1 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0001.01
  Number of active level-1 adjacencies: 1
  Level-1 LSP MTU: 1492
  Next IS-IS LAN Level-1 Hello in 1 seconds

```

```

PE2#show isis interface xe11
xe11 is up, line protocol is up
Routing Protocol: IS-IS (1)
  Network Type: Broadcast
  Circuit Type: level-1
  Local circuit ID: 0x01
  Extended Local circuit ID: 0x0000271B
  Local SNPA: 04f8.f8a1.60f9
  IP interface address:
    10.10.10.2/24
  IPv6 interface address:
    fe80::6f8:f8ff:feal:60f9/64
  Level-1 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0001.01
  Number of active level-1 adjacencies: 1
  Level-1 LSP MTU: 1492

```

Next IS-IS LAN Level-1 Hello in 2 seconds

PE1#show isis database detail

Tag 1: VRF : default

IS-IS Level-1 Link State Database:

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|------------------------|-----------------------------|--------------|--------------|----------|
| 0000.0000.0001.00-00* | 0x00000002 | 0xB193 | 977 | 0/0/0 |
| Area Address: 49.0001 | | | | |
| NLPID: 0xCC | | | | |
| IP Address: 10.10.10.1 | | | | |
| Metric: 10 | IS 0000.0000.0001.01 | | | |
| Metric: 10 | IP 10.10.10.0 255.255.255.0 | | | |
| Metric: 10 | IP 1.1.1.1 255.255.255.255 | | | |
| 0000.0000.0001.01-00* | 0x00000001 | 0x1FBD | 977 | 0/0/0 |
| Metric: 0 | IS 0000.0000.0001.00 | | | |
| Metric: 0 | IS 0000.0000.0002.00 | | | |
| 0000.0000.0002.00-00 | 0x00000002 | 0x84BA | 980 | 0/0/0 |
| Area Address: 49.0001 | | | | |
| NLPID: 0xCC | | | | |
| IP Address: 10.10.10.2 | | | | |
| Metric: 10 | IS 0000.0000.0001.01 | | | |
| Metric: 10 | IP 10.10.10.0 255.255.255.0 | | | |
| Metric: 10 | IP 2.2.2.2 255.255.255.255 | | | |

PE2#show isis database detail

Tag 1: VRF : default

IS-IS Level-1 Link State Database:


| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|------------------------|-----------------------------|--------------|--------------|----------|
| 0000.0000.0001.00-00 | 0x00000002 | 0xB193 | 954 | 0/0/0 |
| Area Address: 49.0001 | | | | |
| NLPID: 0xCC | | | | |
| IP Address: 10.10.10.1 | | | | |
| Metric: 10 | IS 0000.0000.0001.01 | | | |
| Metric: 10 | IP 10.10.10.0 255.255.255.0 | | | |
| Metric: 10 | IP 1.1.1.1 255.255.255.255 | | | |
| 0000.0000.0001.01-00 | 0x00000001 | 0x1FBD | 954 | 0/0/0 |
| Metric: 0 | IS 0000.0000.0001.00 | | | |
| Metric: 0 | IS 0000.0000.0002.00 | | | |
| 0000.0000.0002.00-00* | 0x00000002 | 0x84BA | 958 | 0/0/0 |
| Area Address: 49.0001 | | | | |
| NLPID: 0xCC | | | | |
| IP Address: 10.10.10.2 | | | | |
| Metric: 10 | IS 0000.0000.0001.01 | | | |
| Metric: 10 | IP 10.10.10.0 255.255.255.0 | | | |
| Metric: 10 | IP 2.2.2.2 255.255.255.255 | | | |

LDP-IGP SYNC Configuration

Now that NSM, IS-IS and LDP are all enabled, the LDP-IGP synchronization can be configured.

PE1

| | |
|--|---|
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#mpls ldp-igp sync isis level-1 holddown-timer 700 | Configure LDP-IGP Synchronization for interface xe11 belonging to an IS-IS process with corresponding IS-IS level. 700 seconds is the holddown-timer value for IGP to wait until LDP converges. |


| | |
|---|--|
| | <p>The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command.</p> <p>IS-IS: This command is part of ISIS Process.</p> <p>Default: Mandatory configuration. No default option.</p> <div style="border: 1px solid black; padding: 10px; margin-top: 10px;">  Note: The holddown-timer Range is 1 to 2147483 seconds. If no holddown timer is configured, IGP waits indefinitely for LDP to Converge. Use the command <code>mpls ldp-igp sync is-is <level-type></code> to configure without a holddown-timer. </div> |
| <code>(config-if)#mpls ldp-igp sync-delay 55</code> | <p>Set the time delay in seconds for the notification of LDP convergence to IGP. This is not applicable for notification of non-convergence. Range is 5 to 60 seconds. This command is optional.</p> <p>LDP: This command is part of LDP Process.</p> <p>Default: If not configured, the delay is 0 seconds.</p> |
| <code>(config-if)#commit</code> | Commit the transaction. |

LDP-IGP SYNC Configuration

Now that NSM, IS-IS and LDP are all enabled, the LDP-IGP synchronization can be configured.

PE2

| | |
|--|--|
| <code>(config)#interface xe11</code> | Enter interface mode. |
| <pre>(config-if)#mpls ldp-igp sync isis level-1 holddown-timer 700</pre> | <p>Configure LDP-IGP Synchronization for interface xe11 belonging to an IS-IS process with corresponding IS-IS level. 700 secs is the holddown-timer value for IGP to wait until LDP converges.</p> <p>The parameters level-1 level-2-only level-1-2 identify the IS-IS instance level. The interface can be acting on any level, but sync is applicable only when it matches with the level given in IGP sync command.</p> <p>IS-IS: This command is part of IS-IS Process.</p> <p>Default: Mandatory configuration. No default option.</p> |

| | |
|---|---|
| |  Note: The holddown-timer Range is 1 to 2147483 seconds. If no holddown timer is configured, IGP waits indefinitely for LDP to Converge. Use command <code>mpls ldp-igp sync is-is <level-type></code> to configure without a holddown-timer. |
| <pre>(config-if)#mpls ldp-igp sync-delay 55</pre> | <p>Set the time delay in seconds for notification of LDP convergence to IGP. This is not applicable for notification of non-convergence. Range is 5 to 60 seconds. This command is optional.</p> <p>LDP: This command is part of LDP Process.</p> <p>Default: If not configured, the delay is 0 seconds.</p> |
| <pre>(config-if)#commit</pre> | Commit the transaction. |

PE1 Validation

When LDP IGP SYNC is Configured with hold-down and sync-delay timer

```
PE1#show isis interface xell
xell is up, line protocol is up
  Routing Protocol: IS-IS (1)
    Network Type: Broadcast
    Circuit Type: level-1
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x00002730
    Local SNPA: e8c5.7ad5.3d7a
    IP interface address:
      10.10.10.1/24
    IPv6 interface address:
      fe80::eac5:7aff:fed5:3d7a/64
    LDP-ISIS Sync Configured
      Holddown timer = 700 seconds, Remaining time = 0 seconds
    Level-1 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0001.01
    Number of active level-1 adjacencies: 1
    Level-1 LSP MTU: 1492
    Next IS-IS LAN Level-1 Hello in 394 milliseconds

PE1#show mpls ldp igp sync
Interface      Link-State LDP-Enabled IGP-Sync  Peer-IP      Session-State Sync-Delay-
Time/Remaining-Delay-Time
xell           Up         Yes         Enabled   2.2.2.2      Achieved      55 sec / Not-Running
```

PE2 Validation

```
PE2#show isis interface xell
xell is up, line protocol is up
  Routing Protocol: IS-IS (1)
    Network Type: Broadcast
    Circuit Type: level-1
    Local circuit ID: 0x01
    Extended Local circuit ID: 0x0000271B
    Local SNPA: 04f8.f8a1.60f9
    IP interface address:
      10.10.10.2/24
    IPv6 interface address:
```

```
fe80::6f8:f8ff:feaf:60f9/64
LDP-ISIS Sync Configured
  Holddown timer = 700 seconds, Remaining time = 0 seconds
Level-1 Metric: 10/10, Priority: 64, Circuit ID: 0000.0000.0001.01
Number of active level-1 adjacencies: 1
Level-1 LSP MTU: 1492
Next IS-IS LAN Level-1 Hello in 5 seconds
```

```
PE2#show mpls ldp igp sync
```

| Interface | Link-State | LDP-Enabled | IGP-Sync | Peer-IP | Session-State | Sync-Delay- |
|---------------------------|------------|-------------|----------|---------|---------------|----------------------|
| Time/Remaining-Delay-Time | | | | | | |
| xe11 | Up | Yes | Enabled | 1.1.1.1 | Achieved | 55 sec / Not-Running |

LDP-FRR Configuration

LDP Fast Re-route (FRR) is a technology which helps the router to reduce the MPLS traffic loss in cases of convergence during network failure. A router's convergence time is in general in the order of hundreds of milliseconds, but some applications may be very sensitive to data-loss. This technology helps the router to minimize the MPLS traffic loss by calculating and installing alternate backup paths prior to failure.

LDP FRR improves convergence in case of a single link or single node failure in the network. Convergence times will be in the order of 10s of milliseconds (Max convergence - 50 milliseconds). This is important to some application services (like VoIP) which are sensitive to traffic loss when running over the MPLS network.

Without FRR, link and/or node failures inside an MPLS LDP network result in traffic loss in the order of 100s of milliseconds. The reason for that is that LDP depends on the convergence of the underlying IGP (IS-IS sending LSPs/ OSPF sending LSAs in this case). Post IGP convergence, LDP itself needs to compute new primary Next-Hop Label Forwarding Entries (NHLFEs) for all affected Forwarding Equivalence Classes (FECs). Finally, the different Label Forwarding Information Bases (LFIBs) are updated.

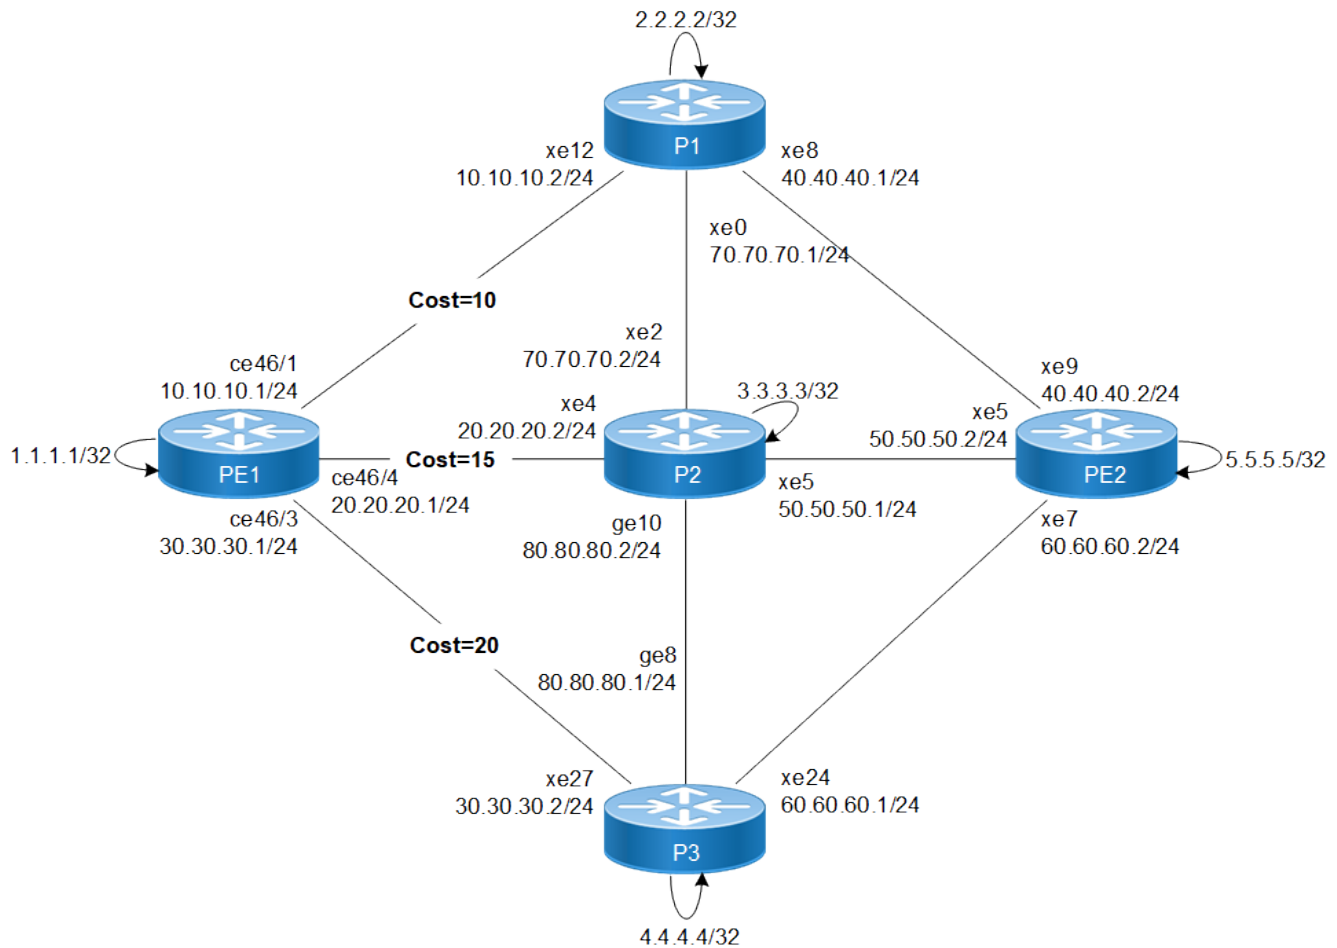
When FRR is configured on a node, the node pre-computes primary NHLFEs for all FECs and in addition it will pre-compute backup NHLFEs for all FECs. The backup NHLFE corresponds to the label received for the same FEC from a Loop-Free Alternate (LFA) next-hop.



Note: This implementation requires either ISIS LFA or OSPF LFA and LDP IGP synchronization.

Topology

Figure 9. LDP-FRR Topology







LDP-FRR with OSPF as IGP Configuration

Below are the configurations and validations involving NSM, OSPF, LDP before configuring fast-reroute for IGP and LDP

PE1


| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1 |



| | |
|---|--|
| (config-router)# fast-reroute | Enables Loop-Free Alternate (LFA) or Remote LFA for rapid failure recovery. |
| (config-router)# session-protection | Enables TCP MD5-based session protection (if configured) and control-plane protection against session flaps. |
| (config-router)#transport-address ipv4 1.1.1.1 0 | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run.  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#targeted-peer ipv4 5.5.5.5 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |
| (config)#interface ce46/1 | Enter interface mode. |
| (config-if)#ip address 10.10.10.1/24 | Configure IPv4 address for xe1. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe1. |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#mpls ldp-igp sync ospf holddown-timer 500 | Enable LDP-IGP Synchronization for xe1 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges. OSPF: This command is part of OSPF Process.  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command mpls ldp-igp sync ospf to configure without a holddown-timer. |
| (config-if)#exit | Exit interface mode |
| (config)#interface ce46/4 | Enter interface mode. |
| (config-if)#ip address 20.20.20.1/24 | Configure IPv4 address for xe2 |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe2. |
| (config-if)#ip ospf cost 15 | Assign ospf cost to the interface |
| (config-if)#mpls ldp-igp sync ospf holddown-timer 500 | Enable LDP-IGP Synchronization for xe1 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges. OSPF: This command is part of OSPF Process. |

| | |
|--|---|
| |  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer. |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface ce46/3</code> | Enter interface mode. |
| <code>(config-if)#ip address 30.30.30.1/24</code> | Configure IPv4 address for xe3 |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe3 |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe3 |
| <code>(config-if)#ip ospf cost 20</code> | Assign OSPF cost to the interface |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for xe1 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p>  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync OSPF</code> to configure without a holddown-timer. |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#router ospf 1</code> | Configure the routing process and specify the Process ID 100. The Process ID should be a unique positive integer identifying the routing process. |
| <code>(config)#ospf router-id 1.1.1.1</code> | Configure OSPF router-ID same as loopback interface IP address |
| <code>(config-router)#network 1.1.1.1/32 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#network 10.10.10.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#network 20.20.20.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#network 30.30.30.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#bfd all-interfaces</code> | Enable the OSPF enabled interfaces with bfd |
| <code>(config-router)#fast-reroute keep-all-paths</code> | Keeps all backup paths in RIB/FIB for use with per- |

| | |
|--|--|
| | prefix FRR. |
| (config-router)# fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp | Enables remote LFA using LDP-based tunnels as repair paths for area 0. |
| (config-if)#exit | Exit interface mode. |
| (config)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval |
| (config)#commit | Commit all the configurations |

P1




| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32 |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| (config-router)#transport-address ipv4 2.2.2.2 0 | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. <div>  Note: It is preferable to use the loopback address as the transport address. </div> |
| (config-router)# fast-reroute | Enables Loop-Free Alternate (LFA) or Remote LFA for rapid failure recovery. |
| (config-router)# session-protection | Enables TCP MD5-based session protection (if configured) and control-plane protection against session flaps. |
| (config-router)#exit | Exit router mode |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 10.10.10.2/24 | Configure IPv4 address for xe1. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe1. |
| (config-if)#mpls ldp-igp sync ospf holddown-timer 500 | Enable LDP-IGP Synchronization for xe1 belonging to an OSPF process. 500 seconds is holddown-timer value for IGP to wait until LDP converges. OSPF: This command is part of OSPF Process. |

| | |
|--|---|
| |  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer. |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface xe8</code> | Enter interface mode. |
| <code>(config-if)#ip address 40.40.40.1/24</code> | Configure IPv4 address for xe2 |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe2. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe2. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for xe1 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p>  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer. |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface xe0</code> | Enter interface mode. |
| <code>(config-if)#ip address 70.70.70.1/24</code> | Configure IPv4 address for xe0. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe0. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe0. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for xe0 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> <p>Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer.</p> |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router ospf 1</code> | Configure the routing process and specify the Process ID |

| | |
|---|---|
| | <ul style="list-style-type: none"> The Process ID should be a unique positive integer identifying the routing process. |
| <code>(config)#ospf router-id 2.2.2.2</code> | Configure OSPF router-ID same as loopback interface IP address |
| <code>(config-router)#network 2.2.2.2/32 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#network 10.10.10.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#network 40.40.40.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#network 70.70.70.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#bfd all-interfaces</code> | Enable the OSPF enabled interfaces with bfd |
| <code>(config-router)#fast-reroute keep-all-paths</code> | Keeps all backup paths in RIB/FIB for use with per-prefix FRR. |
| <code>(config-router)# fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp</code> | Enables remote LFA using LDP-based tunnels as repair paths for area 0. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#bfd interval 3 minrx 3 multiplier 3</code> | Configure BFD interval |
| <code>(config)#commit</code> | Commit all the configurations |

P2

| | |
|---|--|
| <code>#configure terminal</code> | Enter configuration mode. |
| <code>(config)#interface lo</code> | Specify the loopback (lo) interface to be configured. |
| <code>(config-if)#ip address 3.3.3.3/32 secondary</code> | Set the IP address of the loopback interface |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router ldp</code> | Enter router mode for LDP. |
| <code>(config-router)#router-id 3.3.3.3</code> | Set the router ID to IP address 3.3.3.3 |
| <code>(config-router)# fast-reroute</code> | Enables Loop-Free Alternate (LFA) or Remote LFA for rapid failure recovery. |
| <code>(config-router)# session-protection</code> | Enables TCP MD5-based session protection (if configured) and control-plane protection against session flaps. |
| <code>(config-router)#transport-address ipv4 3.3.3.3 0</code> | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. |



| | |
|--|--|
| |  Note: It is preferable to use the loopback address as the transport address. |
| <code>(config-router)#exit</code> | Exit router mode |
| <code>(config)#interface xe4</code> | Enter interface mode. |
| <code>(config-if)#ip address 20.20.20.2/24</code> | Configure IPv4 address for xe4. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe4. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe4. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for xe4 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> <div>  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer. </div> |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface xe5</code> | Enter interface mode. |
| <code>(config-if)#ip address 50.50.50.1/24</code> | Configure IPv4 address for xe5. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe5. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe5. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for xe5 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> <div>  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer. </div> |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface ge10</code> | Enter interface mode. |
| <code>(config-if)#ip address 80.80.80.2/24</code> | Configure IPv4 address for ge10. |

| | |
|--|---|
| <code>(config-if)#label-switching</code> | Enable label switching on interface ge10. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on ge0. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for ge0 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> <p>Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer.</p> |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe2</code> | Enter interface mode. |
| <code>(config-if)#ip address 70.70.70.2/24</code> | Configure IPv4 address for xe2. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe2. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe2. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for xe2 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> <p>Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer.</p> |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router ospf 1</code> | <p>Configure the routing process and specify the Process ID</p> <ul style="list-style-type: none"> The Process ID should be a unique positive integer identifying the routing process. |
| <code>(config)#ospf router-id 3.3.3.3</code> | Configure OSPF router-ID same as loopback interface IP address |
| <code>(config-router)#network 3.3.3.3/32 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#network 20.20.20.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#network 50.50.50.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#network 80.80.80.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#network 70.70.70.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |

| | |
|---|--|
| <code>(config-router)#bfd all-interfaces</code> | Enable the OSPF enabled interfaces with bfd |
| <code>(config-router)#fast-reroute keep-all-paths</code> | Keeps all backup paths in RIB/FIB for use with per-prefix FRR. |
| <code>(config-router)# fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp</code> | Enables remote LFA using LDP-based tunnels as repair paths for area 0. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#bfd interval 3 minrx 3 multiplier 3</code> | Configure BFD interval |
| <code>(config)#commit</code> | Commit all the configurations |

P3


| | |
|--|---|
| <code>#configure terminal</code> | Enter configuration mode. |
| <code>(config)#interface lo</code> | Specify the loopback (lo) interface to be configured. |
| <code>(config-if)#ip address 4.4.4.4/32 secondary</code> | Set the IP address of the loopback interface |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router ldp</code> | Enter router mode for LDP. |
| <code>(config-router)#router-id 4.4.4.4</code> | Set the router ID to IP address 4.4.4.4 |
| <code>(config-router)# fast-reroute</code> | Enables Loop-Free Alternate (LFA) or Remote LFA for rapid failure recovery. |
| <code>(config-router)# session-protection</code> | Enables TCP MD5-based session protection (if configured) and control-plane protection against session flaps. |
| <code>(config-router)#transport-address ipv4 4.4.4.4 0</code> | <p>Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run.</p> <div data-bbox="865 1255 938 1318"> </div> <p>Note: It is preferable to use the loopback address as the transport address.</p> |
| <code>(config-router)#exit</code> | Exit router mode |
| <code>(config)#interface xe27</code> | Enter interface mode. |
| <code>(config-if)#ip address 30.30.30.2/24</code> | Configure IPv4 address for x27. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe27. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe27. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for xe27 belonging to an OSPF process. 500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> |



| | |
|--|---|
| |  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer. |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface ge8</code> | Enter interface mode. |
| <code>(config-if)#ip address 80.80.80.1/24</code> | Configure IPv4 address for ge8. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface ge8. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe0. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for ge8 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> <p>Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer.</p> |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface xe4</code> | Enter interface mode. |
| <code>(config-if)#ip address 60.60.60.1/24</code> | Configure IPv4 address for xe4. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe4. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe4. |
| <code>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</code> | <p>Enable LDP-IGP Synchronization for xe4 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> |
| |  Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer. |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#router ospf 1</code> | <p>Configure the routing process and specify the Process ID</p> <ul style="list-style-type: none"> The Process ID should be a unique positive |


| | |
|---|---|
| | integer identifying the routing process. |
| <code>(config)#ospf router-id 4.4.4.4</code> | Configure OSPF router-ID same as loopback interface IP address. |
| <code>(config-router)#network 4.4.4.4/32 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#network 30.30.30.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#network 60.60.60.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#network 80.80.80.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#bfd all-interfaces</code> | Enable the OSPF enabled interfaces with bfd |
| <code>(config-router)#fast-reroute keep-all-paths</code> | Keeps all backup paths in RIB/FIB for use with per-prefix FRR. |
| <code>(config-router)# fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp</code> | Enables remote LFA using LDP-based tunnels as repair paths for area 0. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#bfd interval 3 minrx 3 multiplier 3</code> | Configure BFD interval |
| <code>(config)#commit</code> | Commit all the configurations |

PE2

| | |
|---|--|
| <code>#configure terminal</code> | Enter configuration mode. |
| <code>(config)#interface lo</code> | Specify the loopback (lo) interface to be configured. |
| <code>(config-if)#ip address 5.5.5.5/32 secondary</code> | Set the IP address of the loopback interface to 5.5.5.5/32 |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router ldp</code> | Enter router mode for LDP. |
| <code>(config-router)#router-id 5.5.5.5</code> | Set the router ID to IP address 5.5.5.5 |
| <code>(config-router)# fast-reroute</code> | Enables Loop-Free Alternate (LFA) or Remote LFA for rapid failure recovery |
| <code>(config-router)# session-protection</code> | Enables TCP MD5-based session protection (if configured) and control-plane protection against session flaps. |
| <code>(config-router)#transport-address ipv4 5.5.5.5 0</code> | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. |


Note: It is preferable to use the loopback address as the transport address.

| | |
|---|--|
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |
| (config)#interface xe9 | Enter interface mode. |
| (config-if)#ip address 40.40.40.2/24 | Configure IPv4 address for xe9. |
| (config-if)#label-switching | Enable label switching on interface xe9. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe9. |
| (config-if)#mpls ldp-igp sync ospf holddown-timer 500 | <p>Enable LDP-IGP Synchronization for xe9 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> <div>  <p>Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command mpls ldp-igp sync ospf to configure without a holddown-timer.</p> </div> |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter interface mode. |
| (config-if)#ip address 50.50.50.2/24 | Configure IPv4 address for xe5. |
| (config-if)#label-switching | Enable label switching on interface xe5. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe5. |
| (config-if)#mpls ldp-igp sync ospf holddown-timer 500 | <p>Enable LDP-IGP Synchronization for xe5 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> <div>  <p>Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command mpls ldp-igp sync ospf to configure without a holddown-timer.</p> </div> |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe7 | Enter interface mode. |
| (config-if)#ip address 60.60.60.2/24 | Configure IPv4 address for xe7. |
| (config-if)#label-switching | Enable label switching on interface xe7. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe7. |

| | |
|---|---|
| <pre>(config-if)#mpls ldp-igp sync ospf holddown-timer 500</pre> | <p>Enable LDP-IGP Synchronization for xe7 belonging to an OSPF process.500 seconds is holddown-timer value for IGP to wait until LDP converges.</p> <p>OSPF: This command is part of OSPF Process.</p> <div>  <p>Note: Holddown-timer range is 1 to 2147483 seconds. If holddown timer is not configured, IGP waits indefinitely for LDP to converge. Use the command <code>mpls ldp-igp sync ospf</code> to configure without a holddown-timer.</p> </div> |
| <pre>(config-if)#exit</pre> | Exit interface mode |
| <pre>(config)#router ospf 1</pre> | <p>Configure the routing process and specify the Process ID</p> <ul style="list-style-type: none"> The Process ID should be a unique positive integer identifying the routing process. |
| <pre>(config)#ospf router-id 5.5.5.5</pre> | Configure OSPF router-ID same as loopback interface IP address |
| <pre>(config-router)#network 5.5.5.5/32 area 0</pre> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <pre>(config-router)#network 40.40.40.0/24 area 0</pre> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <pre>(config-router)#network 50.50.50.0/24 area 0</pre> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <pre>(config-router)#network 60.60.60.0/24 area 0</pre> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <pre>(config-router)#bfd all-interfaces</pre> | Enable the OSPF enabled interfaces with bfd |
| <pre>(config-router)#fast-reroute keep-all-paths</pre> | Keeps all backup paths in RIB/FIB for use with per-prefix FRR. |
| <pre>(config-router)# fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp</pre> | Enables remote LFA using LDP-based tunnels as repair paths for area 0. |
| <pre>(config-if)#exit</pre> | Exit interface mode. |
| <pre>(config)#bfd interval 3 minrx 3 multiplier 3</pre> | Configure BFD interval |
| <pre>(config)#commit</pre> | Commit all the configurations |

Validation

PE1

```
PE1#show ip ospf neighbor
Total number of full neighbors: 3
OSPF process 1 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface      Instance ID
2.2.2.2        1     Full/DR         00:00:31   10.10.10.2    ce46/1         0
3.3.3.3        1     Full/DR         00:00:33   20.20.20.2    ce46/4         0
```

```

4.4.4.4      1      Full/DR      00:00:33      30.30.30.2      ce46/3      0
PE1#show ip ospf interface brief
Interface    PID      Area      Intf ID      Cost      State      Neighbors      Status
lo           1        0.0.0.0    1            1        Loopback      0              Up

ce46/1       1        0.0.0.0    10089        10        Backup      1              Up

ce46/3       1        0.0.0.0    10091        20        Backup      1              Up

ce46/4       1        0.0.0.0    10092        15        Backup      1              Up

PE1#show ip ospf database

        OSPF Router with ID (1.1.1.1) (Process ID 1 VRF default)

        Router Link States (Area 0.0.0.0)

Link ID      ADV Router      Age      Seq#           CkSum      Link count
1.1.1.1      1.1.1.1         870      0x80000016    0xb291     4
2.2.2.2      2.2.2.2         1793     0x80000010    0x9667     4
3.3.3.3      3.3.3.3         890      0x8000001d    0x7402     5
4.4.4.4      4.4.4.4         1698     0x80000016    0x357e     4
5.5.5.5      5.5.5.5         1757     0x80000015    0x1b03     4

        Net Link States (Area 0.0.0.0)

Link ID      ADV Router      Age      Seq#           CkSum
10.10.10.2   2.2.2.2         443      0x8000000c    0x46b4
20.20.20.2   3.3.3.3         600      0x8000000b    0xe2f2
30.30.30.2   4.4.4.4         148      0x8000000c    0x7b33
40.40.40.2   5.5.5.5         487      0x8000000c    0x483c
50.50.50.2   5.5.5.5         347      0x8000000c    0x1151
60.60.60.2   5.5.5.5         1317     0x8000000b    0xdb65
70.70.70.2   3.3.3.3         1340     0x8000000b    0x0635
80.80.80.1   4.4.4.4         1468     0x8000000b    0xdc35

        Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID      ADV Router      Age      Seq#           CkSum      Opaque ID
1.0.0.1      1.1.1.1         360      0x8000000c    0x1a02     1
1.0.0.1      2.2.2.2         73       0x8000000d    0x1cf6     1
1.0.0.1      3.3.3.3         70       0x8000000d    0x20ea     1
1.0.0.1      4.4.4.4         48       0x8000000d    0x24de     1
1.0.0.1      5.5.5.5         157      0x8000000c    0x2ad1     1
1.0.0.8      2.2.2.2         1173     0x8000000b    0x7143     8
1.0.0.10     3.3.3.3         1450     0x8000000b    0x5d50     10
1.0.0.14     3.3.3.3         1060     0x8000000b    0xf8dd     14
1.0.0.16     3.3.3.3         10       0x8000000d    0xb36b     16
1.0.0.16     5.5.5.5         1817     0x8000000c    0x977f     16
1.0.0.20     5.5.5.5         1767     0x8000000c    0x15c1     20
1.0.0.24     2.2.2.2         53       0x8000000c    0xdd7a     24
1.0.0.24     4.4.4.4         1798     0x8000000d    0x10f9     24
1.0.0.24     5.5.5.5         437      0x8000000c    0xa1a9     24
1.0.0.26     3.3.3.3         1700     0x8000000d    0x38d2     26
1.0.0.32     2.2.2.2         1743     0x8000000c    0xba49     32
1.0.0.56     4.4.4.4         1658     0x8000000c    0xab0c     56
1.0.0.62     4.4.4.4         208      0x8000000d    0x9ac9     62
1.0.0.128    1.1.1.1         1650     0x8000000d    0x4c52     128
1.0.0.132    1.1.1.1         120      0x8000000d    0xec2b     132
1.0.0.134    1.1.1.1         360      0x8000000b    0x78e0     134

PE1#show ip route summary

-----
IP routing table name is Default-IP-Routing-Table(0)
-----
IP routing table maximum-paths      : 8
Total number of IPv4 routes         : 14

```

```

Total number of IPv4 paths      : 14
Pending routes (due to route max reached): 0
Route Source      Networks
connected         5
ospf              9
Total            14
FIB              14

```

```

ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
Total number of IPv4 ECMP paths  : 0

```

LFA Non ECMP statistics

```

Total number of Routes      : 9
Total number of Primary Paths : 9
Total number of Backup Paths : 9

```

PE1#show ip interface brief

'*' - address is assigned by dhcp client

| Interface | IP-Address | Admin-Status | Link-Status |
|-----------|------------|--------------|-------------|
| cd48 | unassigned | up | down |
| cd49 | unassigned | up | down |
| cd50 | unassigned | up | down |
| cd51 | unassigned | up | down |
| cd52 | unassigned | up | down |
| cd53 | unassigned | up | down |
| cd54 | unassigned | up | down |
| cd55 | unassigned | up | down |
| ce1 | unassigned | up | down |
| ce2 | unassigned | up | down |
| ce3 | unassigned | up | down |
| ce4 | unassigned | up | down |
| ce5 | unassigned | up | down |
| ce6 | unassigned | up | down |
| ce7 | unassigned | up | down |
| ce8 | unassigned | up | down |
| ce9 | unassigned | up | down |
| ce10 | unassigned | up | down |
| ce11 | unassigned | up | down |
| ce12 | unassigned | up | down |
| ce13 | unassigned | up | down |
| ce14 | unassigned | up | down |
| ce15 | unassigned | up | down |
| ce16 | unassigned | up | down |
| ce17 | unassigned | up | down |
| ce18 | unassigned | up | down |
| ce19 | unassigned | up | down |
| ce20 | unassigned | up | down |
| ce21 | unassigned | up | down |
| ce22 | unassigned | up | down |
| ce23 | unassigned | up | down |
| ce24 | unassigned | up | down |
| ce25 | unassigned | up | down |
| ce26 | unassigned | up | down |
| ce27 | unassigned | up | down |
| ce28 | unassigned | up | down |
| ce29 | unassigned | up | down |
| ce30 | unassigned | up | down |
| ce31 | unassigned | up | down |
| ce32 | unassigned | up | down |
| ce33 | unassigned | up | down |
| ce34 | unassigned | up | down |
| ce35 | unassigned | up | down |
| ce36 | unassigned | up | down |
| ce37 | unassigned | up | down |

```

ce38          unassigned    up          down
ce39          unassigned    up          down
ce40          unassigned    up          down
ce41          unassigned    up          down
ce42          unassigned    up          down
ce43          unassigned    up          down
ce44          unassigned    up          up
ce45          unassigned    up          down
ce46/1        10.10.10.1    up          up
ce46/2        unassigned    up          up
ce46/3        30.30.30.1    up          up
ce46/4        20.20.20.1    up          up
ce47          unassigned    up          down
eth0          *10.16.119.123  up          up
lo            127.0.0.1      up          up
lo.management 127.0.0.1      up          up
xe0          unassigned    up          down
xe1          unassigned    up          up
xe2          unassigned    up          down
xe3          unassigned    up          down

```

PE1#

PE1#show ldp session

```

Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 5.5.5.5 | ce46/1 | Passive | OPERATIONAL | 30 | 05:29:25 |
| | 2.2.2.2 | ce46/1 | Passive | OPERATIONAL | 30 | 05:29:40 |
| | 4.4.4.4 | ce46/3 | Passive | OPERATIONAL | 30 | 05:29:40 |
| | 3.3.3.3 | ce46/4 | Passive | OPERATIONAL | 30 | 05:29:40 |

PE1#show ip route

```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

```

IP Route Table for VRF "default"

```

C      1.1.1.1/32 is directly connected, lo, installed 05:30:36, last update 05:30:36 ago
O      2.2.2.2/32 [110/11] via 10.10.10.2, ce46/1, installed 05:29:42, last update 05:29:42 ago
O      3.3.3.3/32 [110/12] via 10.10.10.2, ce46/1, installed 05:29:42, last update 05:29:42 ago
O      4.4.4.4/32 [110/13] via 10.10.10.2, ce46/1, installed 05:29:22, last update 05:29:22 ago
O      5.5.5.5/32 [110/12] via 10.10.10.2, ce46/1, installed 05:29:42, last update 05:29:42 ago
C      10.10.10.0/24 is directly connected, ce46/1, installed 05:30:33, last update 05:30:33 ago
C      20.20.20.0/24 is directly connected, ce46/4, installed 05:16:01, last update 05:16:01 ago
C      30.30.30.0/24 is directly connected, ce46/3, installed 05:30:33, last update 05:30:33 ago
O      40.40.40.0/24 [110/11] via 10.10.10.2, ce46/1, installed 05:29:42, last update 05:29:42 ago
O      50.50.50.0/24 [110/12] via 10.10.10.2, ce46/1, installed 05:29:16, last update 05:29:16 ago
O      60.60.60.0/24 [110/12] via 10.10.10.2, ce46/1, installed 05:29:19, last update 05:29:19 ago
O      70.70.70.0/24 [110/11] via 10.10.10.2, ce46/1, installed 05:29:44, last update 05:29:42 ago
O      80.80.80.0/24 [110/12] via 10.10.10.2, ce46/1, installed 05:29:22, last update 05:29:22 ago
C      127.0.0.0/8 is directly connected, lo, installed 05:30:57, last update 05:30:57 ago

```

Gateway of last resort is not set

```
PE1#show ldp routes
```

| Prefix | Addr | Nexthop | Addr | Intf | Backup | Addr | Backup |
|---------------|-------|------------|------------|--------|------------|--------|-------------|
| Intf | Owner | CreateTime | UpdateTime | | | | |
| 1.1.1.1/32 | | 0.0.0.0 | | lo | - | - | |
| | | connected | 05:30:57 | - | | | |
| 2.2.2.2/32 | | 10.10.10.2 | | ce46/1 | 20.20.20.2 | ce46/4 | ospf 05:30: |
| 09 05:15:31 | | | | | | | |
| 3.3.3.3/32 | | 10.10.10.2 | | ce46/1 | 20.20.20.2 | ce46/4 | ospf 05:30: |
| 09 05:15:31 | | | | | | | |
| 4.4.4.4/32 | | 10.10.10.2 | | ce46/1 | 20.20.20.2 | ce46/4 | ospf 05:30: |
| 09 05:15:31 | | | | | | | |
| 5.5.5.5/32 | | 10.10.10.2 | | ce46/1 | 20.20.20.2 | ce46/4 | ospf 05:30: |
| 09 05:15:31 | | | | | | | |
| 10.10.10.0/24 | | 0.0.0.0 | | ce46/1 | - | - | |
| | | connected | 05:30:54 | - | | | |
| 20.20.20.0/24 | | 0.0.0.0 | | ce46/4 | - | - | |
| | | connected | 05:16:22 | - | | | |
| 30.30.30.0/24 | | 0.0.0.0 | | ce46/3 | - | - | |
| | | connected | 05:30:54 | - | | | |
| 40.40.40.0/24 | | 10.10.10.2 | | ce46/1 | 30.30.30.2 | ce46/3 | ospf 05:30: |
| 09 05:15:31 | | | | | | | |
| 50.50.50.0/24 | | 10.10.10.2 | | ce46/1 | 20.20.20.2 | ce46/4 | ospf 05:30: |
| 09 05:15:31 | | | | | | | |
| 60.60.60.0/24 | | 10.10.10.2 | | ce46/1 | 20.20.20.2 | ce46/4 | ospf 05:30: |
| 09 05:15:31 | | | | | | | |
| 70.70.70.0/24 | | 10.10.10.2 | | ce46/1 | 20.20.20.2 | ce46/4 | ospf 05:30: |
| 09 05:15:31 | | | | | | | |
| 80.80.80.0/24 | | 10.10.10.2 | | ce46/1 | 20.20.20.2 | ce46/4 | ospf 05:30: |
| 09 05:15:31 | | | | | | | |

```
PE1#show mpls forwarding-table
```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | FTN-ID | Nhlife-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|---------------|---------|-----------|-----------|------------|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 2.2.2.2/32 | 1 | 129 | - | - | - | - |
| | - | N/A | 05:30:17 | - | - | - | - |
| | Yes | 3 | ce46/1 | No | 10.10.10.2 | - | - |
| | No | 24323 | ce46/4 | No | 20.20.20.2 | - | - |
| L> | 3.3.3.3/32 | 2 | 131 | - | - | - | - |
| | - | N/A | 05:30:17 | - | - | - | - |
| | Yes | 24324 | ce46/1 | No | 10.10.10.2 | - | - |
| | No | 3 | ce46/4 | No | 20.20.20.2 | - | - |
| L> | 4.4.4.4/32 | 8 | 135 | - | - | - | - |
| | - | N/A | 05:30:15 | - | - | - | - |
| | Yes | 24321 | ce46/1 | No | 10.10.10.2 | - | - |
| | No | 24324 | ce46/4 | No | 20.20.20.2 | - | - |
| L> | 5.5.5.5/32 | 3 | 138 | - | - | - | - |
| | - | N/A | 05:30:17 | - | - | - | - |
| | Yes | 24325 | ce46/1 | No | 10.10.10.2 | - | - |
| | No | 24325 | ce46/4 | No | 20.20.20.2 | - | - |
| L> | 40.40.40.0/24 | 4 | 141 | - | - | - | - |
| | - | N/A | 05:30:17 | - | - | - | - |
| | Yes | 3 | ce46/1 | No | 10.10.10.2 | - | - |
| | No | 24326 | ce46/3 | No | 30.30.30.2 | - | - |
| L> | 50.50.50.0/24 | 5 | 143 | - | - | - | - |

```

-          N/A          05:30:17
                        19
      Yes  24326          ce46/1    No    10.10.10.2    -    -
                        4
      No    3            ce46/4    No    20.20.20.2    -    -
L> 60.60.60.0/24      9          147    -    -    -    -
-          N/A          05:30:15
                        22
      Yes  24327          ce46/1    No    10.10.10.2    -    -
                        146
      No    24327          ce46/4    No    20.20.20.2    -    -
L> 70.70.70.0/24      6          150    -    -    -    -
-          N/A          05:30:17
                        1
      Yes   3            ce46/1    No    10.10.10.2    -    -
                        4
      No    3            ce46/4    No    20.20.20.2    -    -
L> 80.80.80.0/24      7          152    -    -    -    -
-          N/A          05:30:17
                        25
      Yes  24328          ce46/1    No    10.10.10.2    -    -
                        4
      No    3            ce46/4    No    20.20.20.2    -    -

PE1#show mpls forwarding-table 5.5.5.5/32
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
        B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
        (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code   FEC           FTN-ID   Nhlfe-ID  Tunnel-ID  Pri   Out-Label  Out-
Intf   ELC           Nexthop   Algo-Num  UpTime
L> 5.5.5.5/32      3          138      -          -    -    -    -
-          N/A          05:30:28
                        9
      Yes  24325          ce46/1    No    10.10.10.2    -    -
                        7
      No   24325          ce46/4    No    20.20.20.2    -    -

PE1#show ldp fec
LSR codes   : E/N - LSR is egress/non-egress for this FEC,
              L - LSR received a label for this FEC,
              P - Primary route, B - LFA Backup route,
              R - Remote LFA Backup route,
              > - LSR will use this route for the FEC

FEC           Code   Session           Out Label   ELC   Nexthop Addr
1.1.1.1/32    NL     5.5.5.5           24320      No    connected
              NL     4.4.4.4           24320      No    connected
              NL     2.2.2.2           24320      No    connected
              NL     3.3.3.3           24320      No    connected
              E >   non-existent      none       No    connected
2.2.2.2/32    NL     5.5.5.5           24321      No    no nexthop
              NL     4.4.4.4           24321      No    no nexthop
              NLB>  3.3.3.3           24323      No    20.20.20.2
              NLP>  2.2.2.2           impl-null  No    10.10.10.2
3.3.3.3/32    NL     5.5.5.5           24322      No    no nexthop
              NL     4.4.4.4           24322      No    no nexthop
              NLP>  2.2.2.2           24324      No    10.10.10.2
              NLB>  3.3.3.3           impl-null  No    20.20.20.2
4.4.4.4/32    NL     5.5.5.5           24323      No    no nexthop
              NL     4.4.4.4           impl-null  No    no nexthop
              NLP>  2.2.2.2           24321      No    10.10.10.2
              NLB>  3.3.3.3           24324      No    20.20.20.2
5.5.5.5/32    NL     5.5.5.5           24329      No    no nexthop
              NL     4.4.4.4           24323      No    no nexthop
              NLP>  2.2.2.2           24325      No    10.10.10.2
              NLB>  3.3.3.3           24325      No    20.20.20.2

```



```

10.10.10.0/24    NL      5.5.5.5      24324    No    connected
                 NL      4.4.4.4      24324    No    connected
                 NL      2.2.2.2      impl-null No    connected
                 NL      3.3.3.3      24321    No    connected
                 E >    non-existent none     No    connected
20.20.20.0/24    NL      5.5.5.5      24325    No    connected
                 NL      4.4.4.4      24325    No    connected
                 NL      2.2.2.2      24322    No    connected
                 E >    non-existent none     No    connected
                 NL      3.3.3.3      impl-null No    connected
30.30.30.0/24    NL      5.5.5.5      24326    No    connected
                 NL      4.4.4.4      impl-null No    connected
                 NL      2.2.2.2      24323    No    connected
                 NL      3.3.3.3      24322    No    connected
                 E >    non-existent none     No    connected
40.40.40.0/24    NL      5.5.5.5      24330    No    no nexthop
                 NLB>   4.4.4.4      24326    No    30.30.30.2
                 NL      3.3.3.3      24326    No    no nexthop
                 NLP>   2.2.2.2      impl-null No    10.10.10.2
50.50.50.0/24    NL      5.5.5.5      24331    No    no nexthop
                 NL      4.4.4.4      24327    No    no nexthop
                 NLP>   2.2.2.2      24326    No    10.10.10.2
                 NLB>   3.3.3.3      impl-null No    20.20.20.2
60.60.60.0/24    NL      5.5.5.5      24332    No    no nexthop
                 NL      4.4.4.4      impl-null No    no nexthop
                 NLP>   2.2.2.2      24327    No    10.10.10.2
                 NLB>   3.3.3.3      24327    No    20.20.20.2
70.70.70.0/24    NL      5.5.5.5      24327    No    no nexthop
                 NL      4.4.4.4      24328    No    no nexthop
                 NLP>   2.2.2.2      impl-null No    10.10.10.2
                 NLB>   3.3.3.3      impl-null No    20.20.20.2
80.80.80.0/24    NL      5.5.5.5      24328    No    no nexthop
                 NL      4.4.4.4      impl-null No    no nexthop
                 NLP>   2.2.2.2      24328    No    10.10.10.2
                 NLB>   3.3.3.3      impl-null No    20.20.20.2

```

PE1#show ldp downstream

| FEC | State | Nexthop Label | Req.ID | Attr | Installed | Code |
|--|-----------|-------------------|-------------|-------|-----------|------|
| Codes: P - Primary route, B - Backup route | | | | | | |
| Session peer 5.5.5.5: | | | | | | |
| 20.20.20.0/24 | connected | | Established | 24325 | 0 | None |
| 80.80.80.0/24 | non-nh | Established 24328 | 0 | None | No | - |
| 70.70.70.0/24 | non-nh | Established 24327 | 0 | None | No | - |
| 60.60.60.0/24 | non-nh | Established 24332 | 0 | None | No | - |
| 50.50.50.0/24 | non-nh | Established 24331 | 0 | None | No | - |
| 40.40.40.0/24 | non-nh | Established 24330 | 0 | None | No | - |
| 30.30.30.0/24 | connected | | Established | 24326 | 0 | None |
| 10.10.10.0/24 | connected | | Established | 24324 | 0 | None |
| 5.5.5.5/32 | non-nh | Established 24329 | 0 | None | No | - |
| 4.4.4.4/32 | non-nh | Established 24323 | 0 | None | No | - |
| 3.3.3.3/32 | non-nh | Established 24322 | 0 | None | No | - |
| 2.2.2.2/32 | non-nh | Established 24321 | 0 | None | No | - |
| 1.1.1.1/32 | connected | | Established | 24320 | 0 | None |

[Summary] total downstreams: 13

[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0

```
[label] user-label: 13, impl-null: 0, expl-null: 0, invalid: 0
```

```
Codes: P - Primary route, B - Backup route
```

```
Session peer 2.2.2.2:
```

```
20.20.20.0/24      connected      Established    24322      0      None      No
-
80.80.80.0/24      10.10.10.2    Established    24328      0      None      Ye
s      P
60.60.60.0/24      10.10.10.2    Established    24327      0      None      Ye
s      P
50.50.50.0/24      10.10.10.2    Established    24326      0      None      Ye
s      P
30.30.30.0/24      connected      Established    24323      0      None      No
-
10.10.10.0/24      connected      Established    impl-
null 0      None      No      -
4.4.4.4/32         10.10.10.2    Established    24321      0      None      Ye
s      P
3.3.3.3/32         10.10.10.2    Established    24324      0      None      Ye
s      P
1.1.1.1/32         connected      Established    24320      0      None      No
-
70.70.70.0/24      10.10.10.2    Established    impl-
null 0      None      Yes      P
40.40.40.0/24      10.10.10.2    Established    impl-
null 0      None      Yes      P
5.5.5.5/32         10.10.10.2    Established    24325      0      None      Ye
s      P
2.2.2.2/32         10.10.10.2    Established    impl-
null 0      None      Yes      P
```

```
[Summary] total downstreams: 13
```

```
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
```

```
[label] user-label: 9, impl-null: 4, expl-null: 0, invalid: 0
```

```
Codes: P - Primary route, B - Backup route
```

```
Session peer 4.4.4.4:
```

```
20.20.20.0/24      connected      Established    24325      0      None      No
-
70.70.70.0/24      non-
nh      Established    24328      0      None      No      -
50.50.50.0/24      non-
nh      Established    24327      0      None      No      -
40.40.40.0/24      30.30.30.2    Established    24326      0      None      Ye
s      B
30.30.30.0/24      connected      Established    impl-
null 0      None      No      -
10.10.10.0/24      connected      Established    24324      0      None      No
-
3.3.3.3/32         non-
nh      Established    24322      0      None      No      -
2.2.2.2/32         non-
nh      Established    24321      0      None      No      -
1.1.1.1/32         connected      Established    24320      0      None      No
-
80.80.80.0/24      non-nh
null 0      None      No      -
60.60.60.0/24      non-nh
null 0      None      No      -
5.5.5.5/32         non-
nh      Established    24323      0      None      No      -
4.4.4.4/32         non-nh
null 0      None      No      -
```

```
[Summary] total downstreams: 13
```

```
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
```

```
[label] user-label: 9, impl-null: 4, expl-null: 0, invalid: 0
```

```
Codes: P - Primary route, B - Backup route
```

```
Session peer 3.3.3.3:
```

```
60.60.60.0/24      20.20.20.2    Established    24327      0      None      Ye
s      B
40.40.40.0/24      non-
```

```

nh          Established      24326      0      None      No      -
30.30.30.0/24 connected      Established      24322      0      None      No
-
20.20.20.0/24 connected      Established      impl-
null 0      None      No      -
10.10.10.0/24 connected      Established      24321      0      None      No
-
4.4.4.4/32      20.20.20.2      Established      24324      0      None      Ye
s      B
2.2.2.2/32      20.20.20.2      Established      24323      0      None      Ye
s      B
1.1.1.1/32      connected      Established      24320      0      None      No
-
80.80.80.0/24      20.20.20.2      Established      impl-
null 0      None      Yes      B
70.70.70.0/24      20.20.20.2      Established      impl-
null 0      None      Yes      B
50.50.50.0/24      20.20.20.2      Established      impl-
null 0      None      Yes      B
5.5.5.5/32      20.20.20.2      Established      24325      0      None      Ye
s      B
3.3.3.3/32      20.20.20.2      Established      impl-
null 0      None      Yes      B
[Summary] total downstreams: 13
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
[label] user-label: 8, impl-null: 5, expl-null: 0, invalid: 0

[Summary] total sessions: 4, total downstreams: 52

PE1#show ldp lsp
DOWNSTREAM LSP :
FEC          Nexthop
Addr      State      Label      Req.ID      Attr      Code
1.1.1.1/32 connected      Established      24320      0      None      No
-
1.1.1.1/32 connected      Established      24320      0      None      No
-
1.1.1.1/32 connected      Established      24320      0      None      No
-
1.1.1.1/32 connected      Established      24320      0      None      No
-
1.1.1.1/32 connected      Established      none      0      None      No
-
2.2.2.2/32 non-
nh          Established      24321      0      None      No      -
2.2.2.2/32 non-
nh          Established      24321      0      None      No      -
2.2.2.2/32 20.20.20.2      Established      24323      0      None      Ye
s      B
2.2.2.2/32 10.10.10.2      Established      impl-
null 0      None      Yes      P
3.3.3.3/32 non-
nh          Established      24322      0      None      No      -
3.3.3.3/32 non-
nh          Established      24322      0      None      No      -
3.3.3.3/32 10.10.10.2      Established      24324      0      None      Ye
s      P
3.3.3.3/32 20.20.20.2      Established      impl-
null 0      None      Yes      B
4.4.4.4/32 non-
nh          Established      24323      0      None      No      -
4.4.4.4/32 non-nh      Established      impl-
null 0      None      No      -
4.4.4.4/32 10.10.10.2      Established      24321      0      None      Ye
s      P
4.4.4.4/32 20.20.20.2      Established      24324      0      None      Ye
s      B
5.5.5.5/32 non-
nh          Established      24329      0      None      No      -
5.5.5.5/32 non-

```

| | | | | | | | | |
|---------------|-------------|-------------|-------|------|------|----|----|--|
| nh | Established | 24323 | 0 | None | | No | - | |
| 5.5.5.5/32 | 10.10.10.2 | Established | 24325 | 0 | None | | Ye | |
| s | P | | | | | | | |
| 5.5.5.5/32 | 20.20.20.2 | Established | 24325 | 0 | None | | Ye | |
| s | B | | | | | | | |
| 10.10.10.0/24 | connected | Established | 24324 | 0 | None | | No | |
| - | | | | | | | | |
| 10.10.10.0/24 | connected | Established | 24324 | 0 | None | | No | |
| - | | | | | | | | |
| 10.10.10.0/24 | connected | Established | impl- | | | | | |
| null 0 | None | No | - | | | | | |
| 10.10.10.0/24 | connected | Established | 24321 | 0 | None | | No | |
| - | | | | | | | | |
| 10.10.10.0/24 | connected | Established | none | 0 | None | | No | |
| - | | | | | | | | |
| 20.20.20.0/24 | connected | Established | 24325 | 0 | None | | No | |
| - | | | | | | | | |
| 20.20.20.0/24 | connected | Established | 24325 | 0 | None | | No | |
| - | | | | | | | | |
| 20.20.20.0/24 | connected | Established | 24322 | 0 | None | | No | |
| - | | | | | | | | |
| 20.20.20.0/24 | connected | Established | none | 0 | None | | No | |
| - | | | | | | | | |
| 20.20.20.0/24 | connected | Established | impl- | | | | | |
| null 0 | None | No | - | | | | | |
| 30.30.30.0/24 | connected | Established | 24326 | 0 | None | | No | |
| - | | | | | | | | |
| 30.30.30.0/24 | connected | Established | impl- | | | | | |
| null 0 | None | No | - | | | | | |
| 30.30.30.0/24 | connected | Established | 24323 | 0 | None | | No | |
| - | | | | | | | | |
| 30.30.30.0/24 | connected | Established | 24322 | 0 | None | | No | |
| - | | | | | | | | |
| 30.30.30.0/24 | connected | Established | none | 0 | None | | No | |
| - | | | | | | | | |
| 40.40.40.0/24 | non- | | | | | | | |
| nh | Established | 24330 | 0 | None | | No | - | |
| 40.40.40.0/24 | 30.30.30.2 | Established | 24326 | 0 | None | | Ye | |
| s | B | | | | | | | |
| 40.40.40.0/24 | non- | | | | | | | |
| nh | Established | 24326 | 0 | None | | No | - | |
| 40.40.40.0/24 | 10.10.10.2 | Established | impl- | | | | | |
| null 0 | None | Yes | P | | | | | |
| 50.50.50.0/24 | non- | | | | | | | |
| nh | Established | 24331 | 0 | None | | No | - | |
| 50.50.50.0/24 | non- | | | | | | | |
| nh | Established | 24327 | 0 | None | | No | - | |
| 50.50.50.0/24 | 10.10.10.2 | Established | 24326 | 0 | None | | Ye | |
| s | P | | | | | | | |
| 50.50.50.0/24 | 20.20.20.2 | Established | impl- | | | | | |
| null 0 | None | Yes | B | | | | | |
| 60.60.60.0/24 | non- | | | | | | | |
| nh | Established | 24332 | 0 | None | | No | - | |
| 60.60.60.0/24 | non-nh | | | | | | | |
| null 0 | None | No | - | | | | | |
| 60.60.60.0/24 | 10.10.10.2 | Established | 24327 | 0 | None | | Ye | |
| s | P | | | | | | | |
| 60.60.60.0/24 | 20.20.20.2 | Established | 24327 | 0 | None | | Ye | |
| s | B | | | | | | | |
| 70.70.70.0/24 | non- | | | | | | | |
| nh | Established | 24327 | 0 | None | | No | - | |
| 70.70.70.0/24 | non- | | | | | | | |
| nh | Established | 24328 | 0 | None | | No | - | |
| 70.70.70.0/24 | 10.10.10.2 | Established | impl- | | | | | |
| null 0 | None | Yes | P | | | | | |
| 70.70.70.0/24 | 20.20.20.2 | Established | impl- | | | | | |
| null 0 | None | Yes | B | | | | | |
| 80.80.80.0/24 | non- | | | | | | | |
| nh | Established | 24328 | 0 | None | | No | - | |
| 80.80.80.0/24 | non-nh | | | | | | | |
| null 0 | None | No | - | | | | | |

| | | | | | | |
|----------------|-------------|-------------|--------|------|---------|--------|
| 80.80.80.0/24 | 10.10.10.2 | Established | 24328 | 0 | None | Ye |
| s | | | | | | |
| 80.80.80.0/24 | 20.20.20.2 | Established | impl- | | | |
| null 0 | None | Yes | B | | | |
| UPSTREAM LSP : | | | | | | |
| FEC | State | Label | Req.ID | Attr | | |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes 1 |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes 1 |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes 1 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | skipped | /no 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | skipped | /no 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | skipped | /no 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | skipped | /no 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | skipped | /no 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | skipped | /no 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | skipped | /no 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | skipped | /no 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes 3 |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | skipped | /no 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | skipped | /no 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | skipped | /no 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | skipped | /no 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | skipped | /no 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | skipped | /no 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes 3 |

| | | | | | | | |
|---------------|-------------|-------|---|------|---------|------|---|
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |

PE1#show ldp fec prefix 5.5.5.5/32

LSR codes : E/N - LSR is egress/non-egress for this FEC,
 L - LSR received a label for this FEC,
 P - Primary route, B - LFA Backup route,
 R - Remote LFA Backup route,
 > - LSR will use this route for the FEC

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|------------|------|---------|-----------|-----|--------------|
| 5.5.5.5/32 | NL | 5.5.5.5 | 24329 | No | no nexthop |
| | NL | 4.4.4.4 | 24323 | No | no nexthop |
| | NLP> | 2.2.2.2 | 24325 | No | 10.10.10.2 |
| | NLB> | 3.3.3.3 | 24325 | No | 20.20.20.2 |

PE1#show mpls ftn-table

Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 7, out intf: ce46/1, out label: 3

Nexthop addr: 10.10.10.2 cross connect ix: 3, op code: Push

Backup Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 128 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 128, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/4, out label:

24323

Nexthop addr: 20.20.20.2 cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 2, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 16 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 16, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:

24324

Nexthop addr: 10.10.10.2 cross connect ix: 4, op code: Push

Backup Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3

Nexthop addr: 20.20.20.2 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 8, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 05:33:02, UpTime: 05:33:02, LastUpdate: 05:18:27

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 28 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 28, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:

24321

Nexthop addr: 10.10.10.2 cross connect ix: 9, op code: Push

Backup Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 134 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 134, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/4, out label:

24324

Nexthop addr: 20.20.20.2 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 3, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 9 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 9, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label: 24325

Nexthop addr: 10.10.10.2 cross connect ix: 4, op code: Push

Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: ce46/4, out label: 24325

Nexthop addr: 20.20.20.2 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 40.40.40.0/24, id: 4, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 7, out intf: ce46/1, out label: 3

Nexthop addr: 10.10.10.2 cross connect ix: 3, op code: Push

Backup Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 140 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 140, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/3, out label:

24326

Nexthop addr: 30.30.30.2 cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 50.50.50.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 19 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 19, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:

24326

Nexthop addr: 10.10.10.2 cross connect ix: 6, op code: Push

Backup Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3

Nexthop addr: 20.20.20.2 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 60.60.60.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 05:33:02, UpTime: 05:33:02, LastUpdate: 05:18:27

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 22 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 22, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:

24327

Nexthop addr: 10.10.10.2 cross connect ix: 7, op code: Push

Backup Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 146 bypass ftn-ix: 0

```

Owner: LDP, Persistent: NO, Admin Status: Up, Oper Status: Up
Out-segment with ix: 146, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/4, out label:
24327
  Nexthop addr: 20.20.20.2          cross connect ix: 9, op code: Push

Primary FTN entry with FEC: 70.70.70.0/24, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 refcount: 1
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 7, out intf: ce46/1, out label: 3
    Nexthop addr: 10.10.10.2          cross connect ix: 3, op code: Push

    Backup Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3
    Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 80.80.80.0/24, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 25 refcount: 1
    Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 25, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24328
  Nexthop addr: 10.10.10.2          cross connect ix: 8, op code: Push

    Backup Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3
    Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

PE1#show mpls ftn-table 5.5.5.5/32
Primary FTN entry with FEC: 5.5.5.5/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 05:33:57, UpTime: 05:33:57, LastUpdate: 05:19:20
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 9 refcount: 1
    Owner: LDP, Persistent: NO, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 9, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label: 24325
    Nexthop addr: 10.10.10.2          cross connect ix: 4, op code: Push

    Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 bypass ftn-ix: 0
    Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: ce46/4, out label: 24325
    Nexthop addr: 20.20.20.2          cross connect ix: 4, op code: Push

PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
      K - CLI ILM, T - MPLS-TP, s - Stitched ILM
      S - SNMP, L - LDP, R - RSVP, C - CRLDP
      B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
      O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
      P - SR Policy, U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled
Code  FEC/VRF/L2CKT  ILM-ID  In-Label  Out-Label  In-Intf  Out-
Intf/VRF  Nexthop  pri Algo-Num  UpTime  UPStr peers
L>  4.4.4.4/32      2      24322     24321     N/A      ce46/1    10.10.10.2

```



```

      Yes  N/A      05:34:09  2
L> 2.2.2.2/32      1      24320      3      N/A      ce46/1      10.10.10.2
      Yes  N/A      05:34:09  2
L> 3.3.3.3/32      5      24321      24324      N/A      ce46/1      10.10.10.2
      Yes  N/A      05:34:09  2
L> 50.50.50.0/24   6      24325      24326      N/A      ce46/1      10.10.10.2
      Yes  N/A      05:34:09  2
L> 40.40.40.0/24   3      24324      3      N/A      ce46/1      10.10.10.2
      Yes  N/A      05:34:09  2
                        24324      24326      N/A      ce46/3      30.30.30.2
      No   -      -
L> 5.5.5.5/32      9      24323      24325      N/A      ce46/1      10.10.10.2
      Yes  N/A      05:34:06  2
                        24323      24325      N/A      ce46/4      20.20.20.2
      No   -      -
L> 70.70.70.0/24   8      24327      3      N/A      ce46/1      10.10.10.2
      Yes  N/A      05:34:08  2
                        24327      3      N/A      ce46/4      20.20.20.2
      No   -      -
L> 60.60.60.0/24   4      24326      24327      N/A      ce46/1      10.10.10.2
      Yes  N/A      05:34:09  2
L> 80.80.80.0/24   7      24328      24328      N/A      ce46/1      10.10.10.2
      Yes  N/A      05:34:05  2

PE1#show ldp igp sync
Interface      Link-State LDP-Enabled IGP-Sync  Peer-IP      Session-State Sync-Delay-
Time/Remaining-Delay-Time
ce46/1         Up          Yes         Enabled   2.2.2.2      Achieved      Not-Configured / Not-
Running
ce46/3         Up          Yes         Enabled   4.4.4.4      Achieved      Not-Configured / Not-
Running
ce46/4         Up          Yes         Enabled   3.3.3.3      Achieved      Not-Configured / Not-
Running

PE1#show ldp session multicast count
-----
session up count   : 3
session down count : 0
total count        : 3
-----

PE1#show ldp session targeted count
-----
session up count   : 4
session down count : 0
total count        : 4
-----

PE1#show ldp session all count
-----
session up count   : 4
session down count : 0
total count        : 4
-----

PE1#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity  Adjacency
-----
Multicast 3        3
Targeted  4        4
Total     7        7
-----

PE1#ping 5.5.5.5
Press CTRL+C to exit
PING 5.5.5.5 (5.5.5.5) 100(128) bytes of data.
108 bytes from 5.5.5.5: icmp_seq=1 ttl=64 time=0.602 ms
108 bytes from 5.5.5.5: icmp_seq=2 ttl=64 time=0.508 ms
108 bytes from 5.5.5.5: icmp_seq=3 ttl=64 time=0.624 ms

```

```

108 bytes from 5.5.5.5: icmp_seq=4 ttl=64 time=0.384 ms
108 bytes from 5.5.5.5: icmp_seq=5 ttl=64 time=0.484 ms

--- 5.5.5.5 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4086ms
rtt min/avg/max/mdev = 0.384/0.520/0.624/0.086 ms

```

PE2

```
PE2#show ip ospf neighbor
```

```
Total number of full neighbors: 3
```

```
OSPF process 1 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 2.2.2.2 | 1 | Full/Backup | 00:00:34 | 40.40.40.1 | xe9 | 0 |
| 3.3.3.3 | 1 | Full/Backup | 00:00:34 | 50.50.50.1 | xe5 | 0 |
| 4.4.4.4 | 1 | Full/Backup | 00:00:35 | 60.60.60.1 | xe7 | 0 |

```
PE2#show ip ospf interface brief
```

| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
|-----------|-----|---------|---------|------|----------|-----------|--------|
| lo | 1 | 0.0.0.0 | 1 | 1 | Loopback | 0 | Up |
| xe5 | 1 | 0.0.0.0 | 10005 | 1 | DR | 1 | Up |
| xe7 | 1 | 0.0.0.0 | 10007 | 1 | DR | 1 | Up |
| xe9 | 1 | 0.0.0.0 | 10009 | 1 | DR | 1 | Up |

```
PE2#show ip ospf database
```

```
OSPF Router with ID (5.5.5.5) (Process ID 1 VRF default)
```

```
Router Link States (Area 0.0.0.0)
```

| Link ID | ADV Router | Age | Seq# | CkSum | Link count |
|---------|------------|------|------------|--------|------------|
| 1.1.1.1 | 1.1.1.1 | 352 | 0x80000017 | 0xb092 | 4 |
| 2.2.2.2 | 2.2.2.2 | 1202 | 0x80000011 | 0x9468 | 4 |
| 3.3.3.3 | 3.3.3.3 | 270 | 0x8000001e | 0x7203 | 5 |
| 4.4.4.4 | 4.4.4.4 | 1178 | 0x80000017 | 0x337f | 4 |
| 5.5.5.5 | 5.5.5.5 | 1195 | 0x80000016 | 0x1904 | 4 |

```
Net Link States (Area 0.0.0.0)
```

| Link ID | ADV Router | Age | Seq# | CkSum |
|------------|------------|------|------------|--------|
| 10.10.10.2 | 2.2.2.2 | 1672 | 0x8000000c | 0x46b4 |
| 20.20.20.2 | 3.3.3.3 | 30 | 0x8000000c | 0xe0f3 |
| 30.30.30.2 | 4.4.4.4 | 1378 | 0x8000000c | 0x7b33 |
| 40.40.40.2 | 5.5.5.5 | 1715 | 0x8000000c | 0x483c |
| 50.50.50.2 | 5.5.5.5 | 1575 | 0x8000000c | 0x1151 |
| 60.60.60.2 | 5.5.5.5 | 755 | 0x8000000c | 0xd966 |
| 70.70.70.2 | 3.3.3.3 | 820 | 0x8000000c | 0x0436 |
| 80.80.80.1 | 4.4.4.4 | 848 | 0x8000000c | 0xda36 |

```
Area-Local Opaque-LSA (Area 0.0.0.0)
```

| Link ID | ADV Router | Age | Seq# | CkSum | Opaque ID |
|----------|------------|------|------------|--------|-----------|
| 1.0.0.1 | 1.1.1.1 | 1592 | 0x8000000c | 0x1a02 | 1 |
| 1.0.0.1 | 2.2.2.2 | 1302 | 0x8000000d | 0x1cf6 | 1 |
| 1.0.0.1 | 3.3.3.3 | 1300 | 0x8000000d | 0x20ea | 1 |
| 1.0.0.1 | 4.4.4.4 | 1278 | 0x8000000d | 0x24de | 1 |
| 1.0.0.1 | 5.5.5.5 | 1385 | 0x8000000c | 0x2ad1 | 1 |
| 1.0.0.8 | 2.2.2.2 | 562 | 0x8000000c | 0x6f44 | 8 |
| 1.0.0.10 | 3.3.3.3 | 830 | 0x8000000c | 0x5b51 | 10 |
| 1.0.0.14 | 3.3.3.3 | 500 | 0x8000000c | 0xf6de | 14 |
| 1.0.0.16 | 3.3.3.3 | 1240 | 0x8000000d | 0xb36b | 16 |
| 1.0.0.16 | 5.5.5.5 | 1205 | 0x8000000d | 0x9580 | 16 |
| 1.0.0.20 | 5.5.5.5 | 1195 | 0x8000000d | 0x13c2 | 20 |

| | | | | |
|-----------|---------|------|-------------------|-----|
| 1.0.0.24 | 2.2.2.2 | 1282 | 0x8000000c 0xdd7a | 24 |
| 1.0.0.24 | 4.4.4.4 | 1218 | 0x8000000e 0x0efa | 24 |
| 1.0.0.24 | 5.5.5.5 | 1665 | 0x8000000c 0xa1a9 | 24 |
| 1.0.0.26 | 3.3.3.3 | 1130 | 0x8000000e 0x36d3 | 26 |
| 1.0.0.32 | 2.2.2.2 | 1222 | 0x8000000d 0xb84a | 32 |
| 1.0.0.56 | 4.4.4.4 | 1088 | 0x8000000d 0xa90d | 56 |
| 1.0.0.62 | 4.4.4.4 | 1438 | 0x8000000d 0x9ac9 | 62 |
| 1.0.0.128 | 1.1.1.1 | 1072 | 0x8000000e 0x4a53 | 128 |
| 1.0.0.132 | 1.1.1.1 | 1352 | 0x8000000d 0xec2b | 132 |
| 1.0.0.134 | 1.1.1.1 | 1592 | 0x8000000b 0x78e0 | 134 |

PE2#show ip route summary

IP routing table name is Default-IP-Routing-Table(0)

IP routing table maximum-paths : 8
Total number of IPv4 routes : 14
Total number of IPv4 paths : 18
Pending routes (due to route max reached): 0
Route Source Networks
connected 5
ospf 9
Total 14
FIB 14

ECMP statistics (active in ASIC):

Total number of IPv4 ECMP routes : 3
Total number of IPv4 ECMP paths : 7
Number of routes with 2 ECMP paths: 2
Number of routes with 3 ECMP paths: 1

LFA Non ECMP statistics

Total number of Routes : 4
Total number of Primary Paths : 4
Total number of Backup Paths : 4

LFA ECMP statistics

Total number of Routes : 3
Total number of Primary Paths : 7
Total number of Backup Paths : 7

PE2#show ip interface brief

'*' - address is assigned by dhcp client

| Interface | IP-Address | Admin-Status | Link-Status |
|---------------|----------------|--------------|-------------|
| ce49 | unassigned | up | down |
| ce50 | unassigned | up | down |
| ce51 | unassigned | up | down |
| ce52 | unassigned | up | up |
| ce53 | unassigned | up | down |
| ce54 | unassigned | up | down |
| eth0 | *10.16.119.114 | up | up |
| lo | 127.0.0.1 | up | up |
| lo.management | 127.0.0.1 | up | up |
| xe1 | unassigned | up | up |
| xe2 | unassigned | up | down |
| xe3 | unassigned | up | down |
| xe4 | unassigned | up | up |
| xe5 | 50.50.50.2 | up | up |
| xe6 | unassigned | up | up |
| xe7 | 60.60.60.2 | up | up |
| xe8 | unassigned | up | down |
| xe9 | 40.40.40.2 | up | up |
| xe10 | unassigned | up | down |
| xe11 | unassigned | up | down |

| | | | |
|------|------------|----|------|
| xe12 | unassigned | up | down |
| xe13 | unassigned | up | down |
| xe14 | unassigned | up | down |
| xe15 | unassigned | up | down |
| xe16 | unassigned | up | down |
| xe17 | unassigned | up | down |
| xe18 | unassigned | up | down |
| xe19 | unassigned | up | down |
| xe20 | unassigned | up | down |
| xe21 | unassigned | up | down |
| xe22 | unassigned | up | down |
| xe23 | unassigned | up | down |
| xe24 | unassigned | up | down |
| xe25 | unassigned | up | up |
| xe26 | unassigned | up | down |
| xe27 | unassigned | up | down |
| xe28 | unassigned | up | down |
| xe29 | unassigned | up | down |
| xe30 | unassigned | up | down |
| xe31 | unassigned | up | down |
| xe32 | unassigned | up | down |
| xe33 | unassigned | up | down |
| xe34 | unassigned | up | down |
| xe35 | unassigned | up | down |
| xe36 | unassigned | up | down |
| xe37 | unassigned | up | down |
| xe38 | unassigned | up | down |
| xe39 | unassigned | up | down |
| xe40 | unassigned | up | down |
| xe41 | unassigned | up | down |
| xe42 | unassigned | up | down |
| xe43 | unassigned | up | down |
| xe44 | unassigned | up | down |
| xe45 | unassigned | up | down |
| xe46 | unassigned | up | down |
| xe47 | unassigned | up | down |
| xe48 | unassigned | up | down |

PE2#show ldp session

Codes: m - MD5 password is not set/unset.
 g - GR configuration not set/unset.
 t - TCP MSS not set/unset.
 Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | xe9 | Active | OPERATIONAL | 30 | 05:50:04 |
| | 3.3.3.3 | xe5 | Active | OPERATIONAL | 30 | 05:49:51 |
| | 4.4.4.4 | xe7 | Active | OPERATIONAL | 30 | 05:49:51 |
| | 2.2.2.2 | xe9 | Active | OPERATIONAL | 30 | 05:50:19 |

PE2#show ip route

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
 O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
 ia - IS-IS inter area, E - EVPN,
 v - vrf leaked
 * - candidate default

IP Route Table for VRF "default"

| | | |
|---|---------------|---|
| O | 1.1.1.1/32 | [110/3] via 60.60.60.1, xe7, installed 05:50:40, last update 05:36:08 ago |
| | | [110/3] via 50.50.50.1, xe5 |
| | | [110/3] via 40.40.40.1, xe9 |
| O | 2.2.2.2/32 | [110/2] via 40.40.40.1, xe9, installed 05:50:36, last update 05:50:36 ago |
| O | 3.3.3.3/32 | [110/2] via 50.50.50.1, xe5, installed 05:50:36, last update 05:50:06 ago |
| O | 4.4.4.4/32 | [110/2] via 60.60.60.1, xe7, installed 05:50:11, last update 05:50:11 ago |
| C | 5.5.5.5/32 | is directly connected, lo, installed 05:52:03, last update 05:52:03 ago |
| O | 10.10.10.0/24 | [110/2] via 40.40.40.1, xe9, installed 05:50:36, last update 05:50:36 ago |

```

O      20.20.20.0/24 [110/2] via 50.50.50.1, xe5, installed 05:36:48, last update 05:36:48 ago
O      30.30.30.0/24 [110/2] via 60.60.60.1, xe7, installed 05:50:11, last update 05:50:11 ago
C      40.40.40.0/24 is directly connected, xe9, installed 05:51:27, last update 05:51:27 ago
C      50.50.50.0/24 is directly connected, xe5, installed 05:52:00, last update 05:52:00 ago
C      60.60.60.0/24 is directly connected, xe7, installed 05:52:00, last update 05:52:00 ago
O      70.70.70.0/24 [110/2] via 40.40.40.1, xe9, installed 05:51:11, last update 05:50:06 ago
      [110/2] via 50.50.50.1, xe5
O      80.80.80.0/24 [110/2] via 50.50.50.1, xe5, installed 05:50:11, last update 05:50:06 ago
      [110/2] via 60.60.60.1, xe7
C      127.0.0.0/8 is directly connected, lo, installed 05:52:29, last update 05:52:29 ago

```

Gateway of last resort is not set

PE2#show ldp routes

| Prefix | Addr | Nexthop | Addr | Intf | Backup | Addr | Backup |
|---------------|----------|------------|------------|------|------------------|------|-------------|
| Intf | Owner | CreateTime | UpdateTime | | | | |
| 1.1.1.1/32 | | 40.40.40.1 | | xe9 | 50.50.50.1 | xe5 | ospf 05:51: |
| 00 | 05:36:23 | | | | | | |
| | | 60.60.60.1 | | xe7 | 40.40.40.1 | xe9 | ospf |
| | | 50.50.50.1 | | xe5 | 40.40.40.1 | xe9 | ospf |
| 2.2.2.2/32 | | 40.40.40.1 | | xe9 | 50.50.50.1 | xe5 | ospf 05:51: |
| 02 | 05:36:23 | | | | | | |
| 3.3.3.3/32 | | 50.50.50.1 | | xe5 | 40.40.40.1 | xe9 | ospf 05:51: |
| 31 | 05:36:23 | | | | | | |
| 4.4.4.4/32 | | 60.60.60.1 | | xe7 | 50.50.50.1 | xe5 | ospf 05:51: |
| 26 | 05:36:23 | | | | | | |
| 5.5.5.5/32 | | 0.0.0.0 | | lo | - | - | |
| | | connected | 05:52:23 | - | | | |
| 10.10.10.0/24 | | 40.40.40.1 | | xe9 | 1.1.1.1 | xe5 | ospf 05:51: |
| 02 | 05:50:56 | | | | | | |
| | | | | | (via 50.50.50.1) | | |
| 20.20.20.0/24 | | 50.50.50.1 | | xe5 | 40.40.40.1 | xe9 | ospf 05:37: |
| 13 | 05:36:23 | | | | | | |
| 30.30.30.0/24 | | 60.60.60.1 | | xe7 | 3.3.3.3 | xe9 | ospf 05:51: |
| 26 | 05:36:32 | | | | | | |
| | | | | | (via 40.40.40.1) | | |
| 40.40.40.0/24 | | 0.0.0.0 | | xe9 | - | - | |
| | | connected | 05:51:47 | - | | | |
| 50.50.50.0/24 | | 0.0.0.0 | | xe5 | - | - | |
| | | connected | 05:52:20 | - | | | |
| 60.60.60.0/24 | | 0.0.0.0 | | xe7 | - | - | |
| | | connected | 05:52:20 | - | | | |
| 70.70.70.0/24 | | 40.40.40.1 | | xe9 | 50.50.50.1 | xe5 | ospf 05:51: |
| 31 | 05:36:23 | | | | | | |
| | | 50.50.50.1 | | xe5 | 40.40.40.1 | xe9 | ospf |
| 80.80.80.0/24 | | 60.60.60.1 | | xe7 | 50.50.50.1 | xe5 | ospf 05:51: |
| 31 | 05:36:23 | | | | | | |
| | | 50.50.50.1 | | xe5 | 60.60.60.1 | xe7 | ospf |

PE2#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|------------|---------|----------|-----------|------------|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 1 | 154 | - | - | - | - |
| | - | N/A | 05:51:10 | | | | |
| | | | 1 | - | | | |
| | Yes | 24320 | xe9 | No | 40.40.40.1 | - | - |
| | | | | 28 | - | | |
| | No | 24320 | xe5 | No | 50.50.50.1 | - | - |
| | | | | 22 | - | | |
| | Yes | 24320 | xe7 | No | 60.60.60.1 | - | - |
| | | | | 1 | - | | |
| | No | 24320 | xe9 | No | 40.40.40.1 | - | - |
| | | | | 28 | - | | |

```

      Yes  24320      xe5      No      50.50.50.1      -      -
      No  24320      xe9      No      40.40.40.1      -      -
L> 2.2.2.2/32      2      157      -      -      -      -
      -      N/A      05:51:10
      3      -
      Yes  3      xe9      No      40.40.40.1      -      -
      No  24323      xe5      No      50.50.50.1      -      -
L> 3.3.3.3/32      5      160      -      -      -      -
      -      N/A      05:51:06
      29      -
      Yes  3      xe5      No      50.50.50.1      -      -
      No  24324      xe9      No      40.40.40.1      -      -
L> 4.4.4.4/32      9      166      -      -      -      -
      -      N/A      05:51:05
      24      -
      Yes  3      xe7      No      60.60.60.1      -      -
      No  24324      xe5      No      50.50.50.1      -      -
L> 10.10.10.0/24      3      4      -      -      -      -
      -      N/A      05:51:10
      3      -
      Yes  3      xe9      No      40.40.40.1      -      -
L> 20.20.20.0/24      6      169      -      -      -      -
      -      N/A      05:37:23
      29      -
      Yes  3      xe5      No      50.50.50.1      -      -
      No  24322      xe9      No      40.40.40.1      -      -
L> 30.30.30.0/24      7      153      -      -      -      -
      -      N/A      05:51:06
      24      -
      Yes  3      xe7      No      60.60.60.1      -      -
      No  24322      xe9      No      3.3.3.3      -      -
      (via
40.40.40.1, label 24324)
L> 70.70.70.0/24      4      172      -      -      -      -
      -      N/A      05:51:10
      3      -
      Yes  3      xe9      No      40.40.40.1      -      -
      No  3      xe5      No      50.50.50.1      -      -
      Yes  3      xe5      No      50.50.50.1      -      -
      No  3      xe9      No      40.40.40.1      -      -
L> 80.80.80.0/24      8      174      -      -      -      -
      -      N/A      05:51:06
      24      -
      Yes  3      xe7      No      60.60.60.1      -      -
      No  3      xe5      No      50.50.50.1      -      -
      Yes  3      xe5      No      50.50.50.1      -      -
      No  3      xe7      No      60.60.60.1      -      -

PE2#show mpls forwarding-table 1.1.1.1/32
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code  FEC      FTN-ID  Nhlfe-ID  Tunnel-ID  Pri  Out-Label  Out-
Intf  ELC      Nexthop  Algo-Num  UpTime
L> 1.1.1.1/32      1      154      -      -      -      -

```

```

-                               N/A                               05:51:35
                               1
Yes    24320                   xe9    No    -    40.40.40.1    -    -
                               28
No     24320                   xe5    No    -    50.50.50.1    -    -
                               22
Yes    24320                   xe7    No    -    60.60.60.1    -    -
                               1
No     24320                   xe9    No    -    40.40.40.1    -    -
                               28
Yes    24320                   xe5    No    -    50.50.50.1    -    -
                               1
No     24320                   xe9    No    -    40.40.40.1    -    -

PE2#show ldp routes
Prefix Addr      Nexthop Addr      Intf      Backup Addr      Backup
Intf      Owner      CreateTime  UpdateTime
1.1.1.1/32  40.40.40.1      xe9        50.50.50.1      xe5              ospf      05:51:
46      05:37:09
        60.60.60.1      xe7        40.40.40.1      xe9              ospf
        50.50.50.1      xe5        40.40.40.1      xe9              ospf
2.2.2.2/32  40.40.40.1      xe9        50.50.50.1      xe5              ospf      05:51:
48      05:37:09
3.3.3.3/32  50.50.50.1      xe5        40.40.40.1      xe9              ospf      05:52:
17      05:37:09
4.4.4.4/32  60.60.60.1      xe7        50.50.50.1      xe5              ospf      05:52:
12      05:37:09
5.5.5.5/32  0.0.0.0         lo         -               -
connected  05:53:09      -
10.10.10.0/24 40.40.40.1      xe9        1.1.1.1         xe5              ospf      05:51:
48      05:51:42
        (via 50.50.50.1)
20.20.20.0/24 50.50.50.1      xe5        40.40.40.1      xe9              ospf      05:37:
59      05:37:09
30.30.30.0/24 60.60.60.1      xe7        3.3.3.3         xe9              ospf      05:52:
12      05:37:18
        (via 40.40.40.1)
40.40.40.0/24 0.0.0.0         xe9        -               -
connected  05:52:33      -
50.50.50.0/24 0.0.0.0         xe5        -               -
connected  05:53:06      -
60.60.60.0/24 0.0.0.0         xe7        -               -
connected  05:53:06      -
70.70.70.0/24 40.40.40.1      xe9        50.50.50.1      xe5              ospf      05:52:
17      05:37:09
        50.50.50.1      xe5        40.40.40.1      xe9              ospf
80.80.80.0/24 60.60.60.1      xe7        50.50.50.1      xe5              ospf      05:52:
17      05:37:09
        50.50.50.1      xe5        60.60.60.1      xe7              ospf

PE2#show ldp fec
LSR codes      : E/N - LSR is egress/non-egress for this FEC,
                L - LSR received a label for this FEC,
                P - Primary route, B - LFA Backup route,
                R - Remote LFA Backup route,
                > - LSR will use this route for the FEC
FEC            Code      Session      Out Label      ELC      Nexthop Addr
1.1.1.1/32     NLP>    3.3.3.3     24320          No      50.50.50.1
               NLB>    3.3.3.3     24320          No      50.50.50.1
               NLP>    4.4.4.4     24320          No      60.60.60.1
               NLP>    2.2.2.2     24320          No      40.40.40.1
               NLB>    2.2.2.2     24320          No      40.40.40.1
               NLB>    2.2.2.2     24320          No      40.40.40.1
2.2.2.2/32     NLB>    3.3.3.3     24323          No      50.50.50.1
               NL      4.4.4.4     24321          No      no nexthop
               NLP>    2.2.2.2     impl-null     No      40.40.40.1
3.3.3.3/32     NLP>    3.3.3.3     impl-null     No      50.50.50.1
               NL      4.4.4.4     24322          No      no nexthop
               NLB>    2.2.2.2     24324          No      40.40.40.1

```

```

4.4.4.4/32      NLB> 3.3.3.3      24324      No      50.50.50.1
                NLP> 4.4.4.4      impl-null   No      60.60.60.1
                NL  2.2.2.2      24321      No      no nexthop
5.5.5.5/32      NL  3.3.3.3      24325      No      connected
                NL  4.4.4.4      24323      No      connected
                NL  2.2.2.2      24325      No      connected
                E > non-existent none        No      connected
10.10.10.0/24   NL  3.3.3.3      24321      No      no nexthop
                NL  4.4.4.4      24324      No      no nexthop
                NLP> 2.2.2.2      impl-null   No      40.40.40.1
20.20.20.0/24   NL  4.4.4.4      24325      No      no nexthop
                NLB> 2.2.2.2      24322      No      40.40.40.1
                NLP> 3.3.3.3      impl-null   No      50.50.50.1
30.30.30.0/24   NLR> 3.3.3.3      24322      No      3.3.3.3
                (via 40.40.40.1,label 24324)
                NLP> 4.4.4.4      impl-null   No      60.60.60.1
                NL  2.2.2.2      24323      No      no nexthop
40.40.40.0/24   NL  3.3.3.3      24326      No      connected
                NL  4.4.4.4      24326      No      connected
                NL  2.2.2.2      impl-null   No      connected
                E > non-existent none        No      connected
50.50.50.0/24   NL  3.3.3.3      impl-null   No      connected
                NL  4.4.4.4      24327      No      connected
                NL  2.2.2.2      24326      No      connected
                E > non-existent none        No      connected
60.60.60.0/24   NL  3.3.3.3      24327      No      connected
                NL  4.4.4.4      impl-null   No      connected
                NL  2.2.2.2      24327      No      connected
                E > non-existent none        No      connected
70.70.70.0/24   NLP> 3.3.3.3      impl-null   No      50.50.50.1
                NLB> 3.3.3.3      impl-null   No      50.50.50.1
                NL  4.4.4.4      24328      No      no nexthop
                NLP> 2.2.2.2      impl-null   No      40.40.40.1
                NLB> 2.2.2.2      impl-null   No      40.40.40.1
80.80.80.0/24   NLP> 3.3.3.3      impl-null   No      50.50.50.1
                NLB> 3.3.3.3      impl-null   No      50.50.50.1
                NLP> 4.4.4.4      impl-null   No      60.60.60.1
                NLB> 4.4.4.4      impl-null   No      60.60.60.1
                NL  2.2.2.2      24328      No      no nexthop

```

PE2#show ldp downstream

| FEC | State | Nexthop | Req.ID | Attr | Installed | Code |
|--|-------------|------------|-------------|-------|-----------|------|
| Addr | State | Label | Req.ID | Attr | Installed | Code |
| Codes: P - Primary route, B - Backup route | | | | | | |
| Codes: P - Primary route, B - Backup route | | | | | | |
| Session peer 3.3.3.3: | | | | | | |
| 20.20.20.0/24 | | 50.50.50.1 | Established | impl- | | |
| null 0 | | None | Yes | P | | |
| 80.80.80.0/24 | | 50.50.50.1 | Established | impl- | | |
| null 0 | | None | Yes | P | | |
| 70.70.70.0/24 | | 50.50.50.1 | Established | impl- | | |
| null 0 | | None | Yes | P | | |
| 60.60.60.0/24 | | connected | Established | 24327 | 0 | None |
| | | | | | | No |
| 50.50.50.0/24 | | connected | Established | impl- | | |
| null 0 | | None | No | - | | |
| 40.40.40.0/24 | | connected | Established | 24326 | 0 | None |
| | | | | | | No |
| 30.30.30.0/24 | | 40.40.40.1 | Established | 24322 | 0 | None |
| | | | | | | Ye |
| 10.10.10.0/24 | | non- | | | | |
| nh | Established | 24321 | 0 | None | No | - |
| 5.5.5.5/32 | | connected | Established | 24325 | 0 | None |
| | | | | | | No |
| 4.4.4.4/32 | | 50.50.50.1 | Established | 24324 | 0 | None |
| | | | | | | Ye |
| 3.3.3.3/32 | | 50.50.50.1 | Established | impl- | | |
| null 0 | | None | Yes | P | | |
| 2.2.2.2/32 | | 50.50.50.1 | Established | 24323 | 0 | None |
| | | | | | | Ye |


```

1.1.1.1/32      50.50.50.1      Established      24320      0      None      Ye
s      P
[Summary] total downstreams: 13
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
[label] user-label: 8, impl-null: 5, expl-null: 0, invalid: 0

Codes: P - Primary route, B - Backup route
Session peer 4.4.4.4:
20.20.20.0/24      non-
nh      Established      24325      0      None      No      -
80.80.80.0/24      60.60.60.1      Established      impl-
null 0      None      Yes      P
70.70.70.0/24      non-
nh      Established      24328      0      None      No      -
60.60.60.0/24      connected      Established      impl-
null 0      None      No      -
50.50.50.0/24      connected      Established      24327      0      None      No
-
40.40.40.0/24      connected      Established      24326      0      None      No
-
30.30.30.0/24      60.60.60.1      Established      impl-
null 0      None      Yes      P
10.10.10.0/24      non-
nh      Established      24324      0      None      No      -
5.5.5.5/32      connected      Established      24323      0      None      No
-
4.4.4.4/32      60.60.60.1      Established      impl-
null 0      None      Yes      P
3.3.3.3/32      non-
nh      Established      24322      0      None      No      -
2.2.2.2/32      non-
nh      Established      24321      0      None      No      -
1.1.1.1/32      60.60.60.1      Established      24320      0      None      Ye
s      P
[Summary] total downstreams: 13
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
[label] user-label: 9, impl-null: 4, expl-null: 0, invalid: 0

Codes: P - Primary route, B - Backup route
Session peer 2.2.2.2:
20.20.20.0/24      40.40.40.1      Established      24322      0      None      Ye
s      B
80.80.80.0/24      non-
nh      Established      24328      0      None      No      -
60.60.60.0/24      connected      Established      24327      0      None      No
-
50.50.50.0/24      connected      Established      24326      0      None      No
-
40.40.40.0/24      connected      Established      impl-
null 0      None      No      -
5.5.5.5/32      connected      Established      24325      0      None      No
-
3.3.3.3/32      40.40.40.1      Established      24324      0      None      Ye
s      B
30.30.30.0/24      non-
nh      Established      24323      0      None      No      -
4.4.4.4/32      non-
nh      Established      24321      0      None      No      -
70.70.70.0/24      40.40.40.1      Established      impl-
null 0      None      Yes      P
10.10.10.0/24      40.40.40.1      Established      impl-
null 0      None      Yes      P
2.2.2.2/32      40.40.40.1      Established      impl-
null 0      None      Yes      P
1.1.1.1/32      40.40.40.1      Established      24320      0      None      Ye
s      P
[Summary] total downstreams: 13
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
[label] user-label: 9, impl-null: 4, expl-null: 0, invalid: 0

```

[Summary] total sessions: 3, total downstreams: 39

PE2#show ldp lsp

DOWNSTREAM LSP :

| FEC | State | Nexthop Label | Req.ID | Attr | Code | | | |
|---------------|-------------|---------------|-------------|-------|------|------|--|----|
| 1.1.1.1/32 | P | 50.50.50.1 | Established | 24320 | 0 | None | | Ye |
| 1.1.1.1/32 | P | 60.60.60.1 | Established | 24320 | 0 | None | | Ye |
| 1.1.1.1/32 | P | 40.40.40.1 | Established | 24320 | 0 | None | | Ye |
| 2.2.2.2/32 | B | 50.50.50.1 | Established | 24323 | 0 | None | | Ye |
| 2.2.2.2/32 | Established | non-24321 | 0 | None | No | - | | |
| 2.2.2.2/32 | 0 | 40.40.40.1 | Established | impl- | | | | |
| 3.3.3.3/32 | None | 50.50.50.1 | Yes | P | | | | |
| 3.3.3.3/32 | None | 50.50.50.1 | Yes | P | | | | |
| 3.3.3.3/32 | Established | non-24322 | 0 | None | No | - | | |
| 3.3.3.3/32 | B | 40.40.40.1 | Established | 24324 | 0 | None | | Ye |
| 4.4.4.4/32 | B | 50.50.50.1 | Established | 24324 | 0 | None | | Ye |
| 4.4.4.4/32 | 0 | 60.60.60.1 | Established | impl- | | | | |
| 4.4.4.4/32 | None | 60.60.60.1 | Yes | P | | | | |
| 4.4.4.4/32 | Established | non-24321 | 0 | None | No | - | | |
| 5.5.5.5/32 | - | connected | Established | 24325 | 0 | None | | No |
| 5.5.5.5/32 | - | connected | Established | 24323 | 0 | None | | No |
| 5.5.5.5/32 | - | connected | Established | 24325 | 0 | None | | No |
| 5.5.5.5/32 | - | connected | Established | none | 0 | None | | No |
| 10.10.10.0/24 | Established | non-24321 | 0 | None | No | - | | |
| 10.10.10.0/24 | Established | non-24324 | 0 | None | No | - | | |
| 10.10.10.0/24 | 0 | 40.40.40.1 | Established | impl- | | | | |
| 20.20.20.0/24 | None | 40.40.40.1 | Yes | P | | | | |
| 20.20.20.0/24 | Established | non-24325 | 0 | None | No | - | | |
| 20.20.20.0/24 | B | 40.40.40.1 | Established | 24322 | 0 | None | | Ye |
| 20.20.20.0/24 | 0 | 50.50.50.1 | Established | impl- | | | | |
| 30.30.30.0/24 | None | 40.40.40.1 | Yes | P | | | | |
| 30.30.30.0/24 | B | 40.40.40.1 | Established | 24322 | 0 | None | | Ye |
| 30.30.30.0/24 | 0 | 60.60.60.1 | Established | impl- | | | | |
| 30.30.30.0/24 | None | 60.60.60.1 | Yes | P | | | | |
| 30.30.30.0/24 | Established | non-24323 | 0 | None | No | - | | |
| 40.40.40.0/24 | - | connected | Established | 24326 | 0 | None | | No |
| 40.40.40.0/24 | - | connected | Established | 24326 | 0 | None | | No |
| 40.40.40.0/24 | 0 | connected | Established | impl- | | | | |
| 40.40.40.0/24 | None | connected | No | - | | | | |
| 40.40.40.0/24 | - | connected | Established | none | 0 | None | | No |
| 50.50.50.0/24 | 0 | connected | Established | impl- | | | | |
| 50.50.50.0/24 | None | connected | No | - | | | | |
| 50.50.50.0/24 | - | connected | Established | 24327 | 0 | None | | No |
| 50.50.50.0/24 | - | connected | Established | 24326 | 0 | None | | No |

| | | | | | | |
|----------------|-------------|-------------|--------|------|---------|--------|
| 50.50.50.0/24 | connected | Established | none | 0 | None | No |
| 60.60.60.0/24 | connected | Established | 24327 | 0 | None | No |
| 60.60.60.0/24 | connected | Established | impl- | | | |
| null 0 | None | No | - | | | |
| 60.60.60.0/24 | connected | Established | 24327 | 0 | None | No |
| 60.60.60.0/24 | connected | Established | none | 0 | None | No |
| 70.70.70.0/24 | 50.50.50.1 | Established | impl- | | | |
| null 0 | None | Yes | P | | | |
| 70.70.70.0/24 | non- | | | | | |
| nh | Established | 24328 | 0 | None | No | - |
| 70.70.70.0/24 | 40.40.40.1 | Established | impl- | | | |
| null 0 | None | Yes | P | | | |
| 80.80.80.0/24 | 50.50.50.1 | Established | impl- | | | |
| null 0 | None | Yes | P | | | |
| 80.80.80.0/24 | 60.60.60.1 | Established | impl- | | | |
| null 0 | None | Yes | P | | | |
| 80.80.80.0/24 | non- | | | | | |
| nh | Established | 24328 | 0 | None | No | - |
| UPSTREAM LSP : | | | | | | |
| FEC | State | Label | Req.ID | Attr | | |
| 1.1.1.1/32 | Established | 24320 | 0 | None | Yes | /yes 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | skipped | /no 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | skipped | /no 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | skipped | /no 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | skipped | /no 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | skipped | /no 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | skipped | /no 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | Yes | /yes 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | skipped | /no 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | skipped | /no 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | skipped | /no 4 |
| 1.1.1.1/32 | Established | 24320 | 0 | None | Yes | /yes 4 |
| 2.2.2.2/32 | Established | 24321 | 0 | None | Yes | /yes 4 |
| 2.2.2.2/32 | Established | 24321 | 0 | None | Yes | /yes 4 |
| 2.2.2.2/32 | Established | 24321 | 0 | None | Yes | /yes 4 |
| 2.2.2.2/32 | Established | 24321 | 0 | None | skipped | /no 4 |
| 2.2.2.2/32 | Established | 24321 | 0 | None | skipped | /no 4 |
| 2.2.2.2/32 | Established | 24321 | 0 | None | Yes | /yes 4 |
| 2.2.2.2/32 | Established | 24321 | 0 | None | Yes | /yes 4 |
| 2.2.2.2/32 | Established | 24321 | 0 | None | Yes | /yes 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | skipped | /no 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | Yes | /yes 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | Yes | /yes 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | Yes | /yes 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | Yes | /yes 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | skipped | /no 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | Yes | /yes 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | Yes | /yes 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | Yes | /yes 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | Yes | /yes 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | skipped | /no 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | Yes | /yes 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | Yes | /yes 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | skipped | /no 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | Yes | /yes 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | Yes | /yes 4 |
| 5.5.5.5/32 | Established | impl-null | 0 | None | No | /yes 1 |
| 5.5.5.5/32 | Established | impl-null | 0 | None | No | /yes 1 |
| 5.5.5.5/32 | Established | impl-null | 0 | None | No | /yes 1 |
| 5.5.5.5/32 | Established | 24329 | 0 | None | Yes | /yes 1 |
| 10.10.10.0/24 | Established | 24324 | 0 | None | skipped | /no 4 |
| 10.10.10.0/24 | Established | 24324 | 0 | None | Yes | /yes 4 |
| 10.10.10.0/24 | Established | 24324 | 0 | None | Yes | /yes 4 |
| 10.10.10.0/24 | Established | 24324 | 0 | None | Yes | /yes 4 |
| 20.20.20.0/24 | Established | 24325 | 0 | None | Yes | /yes 4 |

| | | | | | | | |
|---------------|-------------|-----------|---|------|---------|------|---|
| 20.20.20.0/24 | Established | 24325 | 0 | None | skipped | /no | 4 |
| 20.20.20.0/24 | Established | 24325 | 0 | None | Yes | /yes | 4 |
| 20.20.20.0/24 | Established | 24325 | 0 | None | Yes | /yes | 4 |
| 20.20.20.0/24 | Established | 24325 | 0 | None | Yes | /yes | 4 |
| 20.20.20.0/24 | Established | 24325 | 0 | None | skipped | /no | 4 |
| 20.20.20.0/24 | Established | 24325 | 0 | None | Yes | /yes | 4 |
| 20.20.20.0/24 | Established | 24325 | 0 | None | Yes | /yes | 4 |
| 30.30.30.0/24 | Established | 24326 | 0 | None | Yes | /yes | 4 |
| 30.30.30.0/24 | Established | 24326 | 0 | None | Yes | /yes | 4 |
| 30.30.30.0/24 | Established | 24326 | 0 | None | skipped | /no | 4 |
| 30.30.30.0/24 | Established | 24326 | 0 | None | Yes | /yes | 4 |
| 30.30.30.0/24 | Established | 24326 | 0 | None | Yes | /yes | 4 |
| 30.30.30.0/24 | Established | 24326 | 0 | None | skipped | /no | 4 |
| 30.30.30.0/24 | Established | 24326 | 0 | None | Yes | /yes | 4 |
| 30.30.30.0/24 | Established | 24326 | 0 | None | Yes | /yes | 4 |
| 40.40.40.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 40.40.40.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 40.40.40.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 40.40.40.0/24 | Established | 24330 | 0 | None | Yes | /yes | 1 |
| 50.50.50.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 50.50.50.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 50.50.50.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 50.50.50.0/24 | Established | 24331 | 0 | None | Yes | /yes | 1 |
| 60.60.60.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 60.60.60.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 60.60.60.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 60.60.60.0/24 | Established | 24332 | 0 | None | Yes | /yes | 1 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no | 4 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 4 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no | 4 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 4 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no | 4 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 4 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no | 4 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 4 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no | 4 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no | 4 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 4 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 4 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no | 4 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no | 4 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 4 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 4 |

PE2#show ldp fec prefix 1.1.1.1/32

LSR codes : E/N - LSR is egress/non-egress for this FEC,
 L - LSR received a label for this FEC,
 P - Primary route, B - LFA Backup route,
 R - Remote LFA Backup route,
 > - LSR will use this route for the FEC

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|------------|------|---------|-----------|-----|--------------|
| 1.1.1.1/32 | NLP> | 3.3.3.3 | 24320 | No | 50.50.50.1 |
| | NLB> | 3.3.3.3 | 24320 | No | 50.50.50.1 |
| | NLP> | 4.4.4.4 | 24320 | No | 60.60.60.1 |
| | NLP> | 2.2.2.2 | 24320 | No | 40.40.40.1 |
| | NLB> | 2.2.2.2 | 24320 | No | 40.40.40.1 |
| | NLB> | 2.2.2.2 | 24320 | No | 40.40.40.1 |

PE2#show router-id

Name: management
 Router ID: 10.16.119.114 (automatic)
 Name: default
 Router ID: 5.5.5.5 (loopback)

PE2#show mpls ftn-table

Primary FTN entry with FEC: 1.1.1.1/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed
 CreateTime: 05:53:13, UpTime: 05:53:13, LastUpdate: 05:38:36
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,

Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: LDP, Stale: NO, refcount: 7, out intf: xe9, out label: 24320

Nexthop addr: 40.40.40.1 cross connect ix: 2, op code: Push

Backup Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 28 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 28, owner: LDP, Stale: NO, refcount: 6, out intf: xe5, out label: 24320

Nexthop addr: 50.50.50.1 cross connect ix: 2, op code: Push

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 22 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 22, owner: LDP, Stale: NO, refcount: 2, out intf: xe7, out label: 24320

Nexthop addr: 60.60.60.1 cross connect ix: 2, op code: Push

Backup Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: LDP, Stale: NO, refcount: 7, out intf: xe9, out label: 24320

Nexthop addr: 40.40.40.1 cross connect ix: 2, op code: Push

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 28 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 28, owner: LDP, Stale: NO, refcount: 6, out intf: xe5, out label: 24320

Nexthop addr: 50.50.50.1 cross connect ix: 2, op code: Push

Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 1 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: LDP, Stale: NO, refcount: 7, out intf: xe9, out label: 24320

Nexthop addr: 40.40.40.1 cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 2.2.2.2/32, id: 2, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 05:53:13, UpTime: 05:53:13, LastUpdate: 05:38:36

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,

Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 3, owner: N/A, Stale: NO, refcount: 8, out intf: xe9, out label: 3

Nexthop addr: 40.40.40.1 cross connect ix: 3, op code: Push

Backup Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 156 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 156, owner: LDP, Stale: NO, refcount: 2, out intf: xe5, out label: 24323

Nexthop addr: 50.50.50.1 cross connect ix: 8, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 5, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 05:53:09, UpTime: 05:53:09, LastUpdate: 05:38:36

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,

Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 29 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 29, owner: N/A, Stale: NO, refcount: 17, out intf: xe5, out label: 3

Nexthop addr: 50.50.50.1 cross connect ix: 1, op code: Push

Backup Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 5 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 5, owner: LDP, Stale: NO, refcount: 4, out intf: xe9, out label: 24324

Nexthop addr: 40.40.40.1 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 9, row status: Active, Tunnel-Policy: N/A, State: Installed

```
CreateTime: 05:53:08, UpTime: 05:53:08, LastUpdate: 05:38:36
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 24 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 24, owner: N/A, Stale: NO, refcount: 8, out intf: xe7, out label: 3
  Nexthop addr: 60.60.60.1      cross connect ix: 1, op code: Push

  Backup Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 165 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 165, owner: LDP, Stale: NO, refcount: 2, out intf: xe5, out label: 24324
  Nexthop addr: 50.50.50.1      cross connect ix: 13, op code: Push

Primary FTN entry with FEC: 10.10.10.0/24, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 05:53:13, UpTime: 05:53:13, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 3, owner: N/A, Stale: NO, refcount: 8, out intf: xe9, out label: 3
  Nexthop addr: 40.40.40.1      cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 20.20.20.0/24, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 05:39:26, UpTime: 05:39:26, LastUpdate: 05:38:36
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 29 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 29, owner: N/A, Stale: NO, refcount: 17, out intf: xe5, out label: 3
  Nexthop addr: 50.50.50.1      cross connect ix: 1, op code: Push

  Backup Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 7 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: xe9, out label: 24322
  Nexthop addr: 40.40.40.1      cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 30.30.30.0/24, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 05:53:09, UpTime: 05:53:09, LastUpdate: 05:38:36
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 24 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 24, owner: N/A, Stale: NO, refcount: 8, out intf: xe7, out label: 3
  Nexthop addr: 60.60.60.1      cross connect ix: 1, op code: Push

  Backup Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 152 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 152, owner: LDP, Stale: NO, refcount: 2, out intf: xe9, out label: 24322
  Nexthop addr: 3.3.3.3      cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 70.70.70.0/24, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 05:53:13, UpTime: 05:53:13, LastUpdate: 05:38:36
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 3, owner: N/A, Stale: NO, refcount: 8, out intf: xe9, out label: 3
Nexthop addr: 40.40.40.1          cross connect ix: 3, op code: Push

Backup Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 29 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 29, owner: N/A, Stale: NO, refcount: 17, out intf: xe5, out label: 3
Nexthop addr: 50.50.50.1          cross connect ix: 1, op code: Push

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 29 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 29, owner: N/A, Stale: NO, refcount: 17, out intf: xe5, out label: 3
Nexthop addr: 50.50.50.1          cross connect ix: 1, op code: Push

Backup Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 3 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: N/A, Stale: NO, refcount: 8, out intf: xe9, out label: 3
Nexthop addr: 40.40.40.1          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 80.80.80.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 05:53:09, UpTime: 05:53:09, LastUpdate: 05:38:36
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 24 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 24, owner: N/A, Stale: NO, refcount: 8, out intf: xe7, out label: 3
Nexthop addr: 60.60.60.1          cross connect ix: 1, op code: Push

Backup Cross connect ix: 18, in intf: - in label: 0 out-segment ix: 29 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 29, owner: N/A, Stale: NO, refcount: 17, out intf: xe5, out label: 3
Nexthop addr: 50.50.50.1          cross connect ix: 1, op code: Push

Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 29 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 29, owner: N/A, Stale: NO, refcount: 17, out intf: xe5, out label: 3
Nexthop addr: 50.50.50.1          cross connect ix: 1, op code: Push

Backup Cross connect ix: 19, in intf: - in label: 0 out-segment ix: 24 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 24, owner: N/A, Stale: NO, refcount: 8, out intf: xe7, out label: 3
Nexthop addr: 60.60.60.1          cross connect ix: 1, op code: Push

PE2#show mpls ftn-table 1.1.1.1/32
Primary FTN entry with FEC: 1.1.1.1/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 05:53:39, UpTime: 05:53:39, LastUpdate: 05:39:02
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: LDP, Stale: NO, refcount: 7, out intf: xe9, out label: 24320
Nexthop addr: 40.40.40.1          cross connect ix: 2, op code: Push

Backup Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 28 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 28, owner: LDP, Stale: NO, refcount: 6, out intf: xe5, out label: 24320
Nexthop addr: 50.50.50.1          cross connect ix: 2, op code: Push

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 22 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 22, owner: LDP, Stale: NO, refcount: 2, out intf: xe7, out label: 24320
Nexthop addr: 60.60.60.1          cross connect ix: 2, op code: Push
```

```

Backup Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: LDP, Stale: NO, refcount: 7, out intf: xe9, out label: 24320
Nexthop addr: 40.40.40.1 cross connect ix: 2, op code: Push

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 28 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 28, owner: LDP, Stale: NO, refcount: 6, out intf: xe5, out label: 24320
Nexthop addr: 50.50.50.1 cross connect ix: 2, op code: Push

Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 1 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: LDP, Stale: NO, refcount: 7, out intf: xe9, out label: 24320
Nexthop addr: 40.40.40.1 cross connect ix: 2, op code: Push

```

PE2#show mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup

K - CLI ILM, T - MPLS-TP, s - Stitched ILM
 S - SNMP, L - LDP, R - RSVP, C - CRLDP
 B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
 O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
 P - SR Policy, U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out-Intf | UPStr | peers |
|--------------------------|---------------|--------|----------|-----------|---------|----------|-------|------------|
| Intf/VRF | Nexthop | | pri | Algo-Num | UpTime | | | |
| L> 30.30.30.0/24 | Yes N/A | 7 | 24326 | 3 | N/A | xe7 | | 60.60.60.1 |
| | No - | - | 24326 | 24322 | N/A | xe9 | | 3.3.3.3 |
| | | | | | | | | (via |
| 40.40.40.1, label 24324) | | | | | | | | |
| L> 4.4.4.4/32 | Yes N/A | 4 | 24323 | 3 | N/A | xe7 | | 60.60.60.1 |
| | No - | - | 24323 | 24324 | N/A | xe5 | | 50.50.50.1 |
| L> 2.2.2.2/32 | Yes N/A | 2 | 24321 | 3 | N/A | xe9 | | 40.40.40.1 |
| | No - | - | 24321 | 24323 | N/A | xe5 | | 50.50.50.1 |
| L> 1.1.1.1/32 | Yes N/A | 6 | 24320 | 24320 | N/A | xe9 | | 40.40.40.1 |
| | No - | - | 24320 | 24320 | N/A | xe5 | | 50.50.50.1 |
| | Yes - | - | 24320 | 24320 | N/A | xe7 | | 60.60.60.1 |
| | No - | - | 24320 | 24320 | N/A | xe9 | | 40.40.40.1 |
| | Yes - | - | 24320 | 24320 | N/A | xe5 | | 50.50.50.1 |
| | No - | - | 24320 | 24320 | N/A | xe9 | | 40.40.40.1 |
| L> 3.3.3.3/32 | Yes N/A | 1 | 24322 | 3 | N/A | xe5 | | 50.50.50.1 |
| | No - | - | 24322 | 24324 | N/A | xe9 | | 40.40.40.1 |
| L> 10.10.10.0/24 | Yes N/A | 5 | 24324 | 3 | N/A | xe9 | | 40.40.40.1 |
| L> 20.20.20.0/24 | Yes N/A | 3 | 24325 | 3 | N/A | xe5 | | 50.50.50.1 |
| L> 40.40.40.0/24 | Yes N/A | 11 | 24330 | NoLabel | N/A | N/A | | 127.0.0.1 |
| L> 80.80.80.0/24 | Yes N/A | 9 | 24328 | 3 | N/A | xe7 | | 60.60.60.1 |
| | No - | - | 24328 | 3 | N/A | xe5 | | 50.50.50.1 |
| | Yes - | - | 24328 | 3 | N/A | xe5 | | 50.50.50.1 |
| | | | 24328 | 3 | N/A | xe7 | | 60.60.60.1 |


```

      No  -      -
L>  70.70.70.0/24      8      24327      3      N/A      xe9      40.40.40.1
      Yes N/A      05:53:20      2
      24327      3      N/A      xe5      50.50.50.1
      No  -      -
      24327      3      N/A      xe5      50.50.50.1
      Yes -      -
      24327      3      N/A      xe9      40.40.40.1
      No  -      -
L>  5.5.5.5/32      10      24329      Nolabel      N/A      N/A      127.0.0.1
      Yes N/A      05:45:26      1
L>  50.50.50.0/24      12      24331      Nolabel      N/A      N/A      127.0.0.1
      Yes N/A      05:45:26      1
L>  60.60.60.0/24      13      24332      Nolabel      N/A      N/A      127.0.0.1
      Yes N/A      05:45:26      1

PE2# show ldp igp sync
Interface      Link-State LDP-Enabled IGP-Sync Peer-IP Session-State Sync-Delay-
Time/Remaining-Delay-Time
xe5            Up      Yes      Enabled  3.3.3.3      Achieved      Not-Configured / Not-
Running
xe7            Up      Yes      Enabled  4.4.4.4      Achieved      Not-Configured / Not-
Running
xe9            Up      Yes      Enabled  2.2.2.2      Achieved      Not-Configured / Not-
Running
PE2#show ldp
Router ID      : 5.5.5.5
LDP Version    : 1
Fast-reroute Per-prefix : Enabled
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode : Liberal
Label Control Mode : Independent
Instance Loop Detection : Off
Request Retry   : Off
Propagate Release : Disabled
Graceful Restart : Disabled
Hello Interval  : 5
Targeted Hello Interval : 15
Hold time       : 15
Targeted Hold time : 45
Keepalive Interval : 10
Keepalive Timeout : 30
Request retry Timeout : 5
Auto Targeted Hello Receipt : Enabled
Transport Address data :
  Labelspace 0 : 5.5.5.5 (in use)
Import BGP routes : No
Entropy-Label Capability : Disabled
Prefer-Tunnel-In-Tunnel : Disabled

PE2#show ldp session detail
Peer IP Address: 1.1.1.1
=====
Session state      : OPERATIONAL
Session role       : Active
TCP Connection     : Established
IP Address for TCP : 1.1.1.1
Interface being used : xe9
Peer LDP ID        : 1.1.1.1:0
Last Down Reason   : None
Preferred Peer LDP Password : Not Set
Adjacencies        : 1.1.1.1
Advertisement mode  : Downstream Unsolicited
Label retention mode : Liberal
Graceful Restart   : Not Capable
Keepalive Timeout  : 30
Reconnect Interval : 15
Configured TCP MSS : Not configured
Applied TCP MSS    : 1460

```

```

Preferred TCP MSS      : NA
Session protection (SP) : Incomplete
Address List received  : 1.1.1.1
                        10.10.10.1
                        20.20.20.1
                        30.30.30.1

Received Labels :      Fec      Label      Maps To
Sent Labels :   Fec      Label      Maps To
IPV4:20.20.20.0/24  24325      impl-null
IPV4:20.20.20.0/24  24325      24322
IPV4:80.80.80.0/24  24328      impl-null
IPV4:80.80.80.0/24  24328      impl-null
IPV4:70.70.70.0/24  24327      impl-null
IPV4:70.70.70.0/24  24327      impl-null
IPV4:60.60.60.0/24  24332      none
IPV4:50.50.50.0/24  24331      none
IPV4:40.40.40.0/24  24330      none
IPV4:30.30.30.0/24  24326      impl-null
IPV4:30.30.30.0/24  24326      24322
IPV4:10.10.10.0/24  24324      impl-null
IPV4:5.5.5.5/32     24329      none
IPV4:4.4.4.4/32     24323      impl-null
IPV4:4.4.4.4/32     24323      24324
IPV4:3.3.3.3/32     24322      impl-null
IPV4:3.3.3.3/32     24322      24324
IPV4:2.2.2.2/32     24321      impl-null
IPV4:2.2.2.2/32     24321      24323
IPV4:1.1.1.1/32     24320      24320
IPV4:1.1.1.1/32     24320      24320
IPV4:1.1.1.1/32     24320      24320

```

Peer IP Address: 3.3.3.3

```

=====
Session state          : OPERATIONAL
Session role          : Active
TCP Connection         : Established
IP Address for TCP     : 3.3.3.3
Interface being used   : xe5
Peer LDP ID            : 3.3.3.3:0
Last Down Reason       : None
Preferred Peer LDP Password : Not Set
Adjacencies           : 50.50.50.1
                       3.3.3.3
Advertisement mode      : Downstream Unsolicited
Label retention mode    : Liberal
Graceful Restart       : Not Capable
Keepalive Timeout      : 30
Reconnect Interval     : 15
Configured TCP MSS     : Not configured
Applied TCP MSS        : 1460
Preferred TCP MSS      : NA
Session protection (SP) : Ready
SP duration            : infinite
Address List received  : 3.3.3.3
                       20.20.20.2
                       50.50.50.1
                       70.70.70.2
                       80.80.80.2

```

```

Received Labels :      Fec      Label      Maps To
IPV4:20.20.20.0/24  impl-null  24325
IPV4:20.20.20.0/24  impl-null  24325
IPV4:20.20.20.0/24  impl-null  24325
IPV4:20.20.20.0/24  impl-null  24325
IPV4:80.80.80.0/24  impl-null  24328
IPV4:80.80.80.0/24  impl-null  24328
IPV4:80.80.80.0/24  impl-null  24328
IPV4:80.80.80.0/24  impl-null  24328
IPV4:70.70.70.0/24  impl-null  24327

```

| | | |
|--------------------|--------------------|-----------|
| IPv4:70.70.70.0/24 | impl-null | 24327 |
| IPv4:70.70.70.0/24 | impl-null | 24327 |
| IPv4:70.70.70.0/24 | impl-null | 24327 |
| IPv4:60.60.60.0/24 | 24327 | none |
| IPv4:50.50.50.0/24 | impl-null | none |
| IPv4:40.40.40.0/24 | 24326 | none |
| IPv4:30.30.30.0/24 | 24322 | 24326 |
| IPv4:30.30.30.0/24 | 24322 | 24326 |
| IPv4:30.30.30.0/24 | 24322 | 24326 |
| IPv4:30.30.30.0/24 | 24322 | 24326 |
| IPv4:10.10.10.0/24 | 24321 | none |
| IPv4:5.5.5.5/32 | 24325 | none |
| IPv4:4.4.4.4/32 | 24324 | 24323 |
| IPv4:4.4.4.4/32 | 24324 | 24323 |
| IPv4:4.4.4.4/32 | 24324 | 24323 |
| IPv4:4.4.4.4/32 | 24324 | 24323 |
| IPv4:3.3.3.3/32 | impl-null | 24322 |
| IPv4:3.3.3.3/32 | impl-null | 24322 |
| IPv4:3.3.3.3/32 | impl-null | 24322 |
| IPv4:3.3.3.3/32 | impl-null | 24322 |
| IPv4:2.2.2.2/32 | 24323 | 24321 |
| IPv4:2.2.2.2/32 | 24323 | 24321 |
| IPv4:2.2.2.2/32 | 24323 | 24321 |
| IPv4:2.2.2.2/32 | 24323 | 24321 |
| IPv4:1.1.1.1/32 | 24320 | 24320 |
| IPv4:1.1.1.1/32 | 24320 | 24320 |
| IPv4:1.1.1.1/32 | 24320 | 24320 |
| IPv4:1.1.1.1/32 | 24320 | 24320 |
| Sent Labels : | Fec | Label |
| | IPv4:20.20.20.0/24 | 24325 |
| | IPv4:20.20.20.0/24 | 24325 |
| | IPv4:80.80.80.0/24 | 24328 |
| | IPv4:80.80.80.0/24 | 24328 |
| | IPv4:70.70.70.0/24 | 24327 |
| | IPv4:70.70.70.0/24 | 24327 |
| | IPv4:60.60.60.0/24 | impl-null |
| | IPv4:50.50.50.0/24 | impl-null |
| | IPv4:40.40.40.0/24 | impl-null |
| | IPv4:30.30.30.0/24 | 24326 |
| | IPv4:30.30.30.0/24 | 24326 |
| | IPv4:10.10.10.0/24 | 24324 |
| | IPv4:5.5.5.5/32 | impl-null |
| | IPv4:4.4.4.4/32 | 24323 |
| | IPv4:4.4.4.4/32 | 24323 |
| | IPv4:3.3.3.3/32 | 24322 |
| | IPv4:3.3.3.3/32 | 24322 |
| | IPv4:2.2.2.2/32 | 24321 |
| | IPv4:2.2.2.2/32 | 24321 |
| | IPv4:1.1.1.1/32 | 24320 |
| | IPv4:1.1.1.1/32 | 24320 |
| | IPv4:1.1.1.1/32 | 24320 |
| | IPv4:1.1.1.1/32 | 24320 |

Peer IP Address: 4.4.4.4

```

=====
Session state           : OPERATIONAL
Session role           : Active
TCP Connection          : Established
IP Address for TCP      : 4.4.4.4
Interface being used    : xe7
Peer LDP ID             : 4.4.4.4:0
Last Down Reason        : None
Preferred Peer LDP Password : Not Set
Adjacencies             : 60.60.60.1
                        : 4.4.4.4
Advertisement mode       : Downstream Unsolicited
Label retention mode     : Liberal
Graceful Restart        : Not Capable
Keepalive Timeout       : 30

```

```

Reconnect Interval      : 15
Configured TCP MSS     : Not configured
Applied TCP MSS        : 1460
Preferred TCP MSS      : NA
Session protection (SP) : Ready
SP duration            : infinite
Address List received  : 4.4.4.4
                        : 30.30.30.2
                        : 60.60.60.1
                        : 80.80.80.1

Received Labels :      Fec          Label          Maps To
                  IPv4:20.20.20.0/24  24325         none
                  IPv4:80.80.80.0/24  impl-null     24328
                  IPv4:80.80.80.0/24  impl-null     24328
                  IPv4:80.80.80.0/24  impl-null     24328
                  IPv4:80.80.80.0/24  impl-null     24328
                  IPv4:70.70.70.0/24  24328         none
                  IPv4:60.60.60.0/24  impl-null     none
                  IPv4:50.50.50.0/24  24327         none
                  IPv4:40.40.40.0/24  24326         none
                  IPv4:30.30.30.0/24  impl-null     24326
                  IPv4:30.30.30.0/24  impl-null     24326
                  IPv4:30.30.30.0/24  impl-null     24326
                  IPv4:30.30.30.0/24  impl-null     24326
                  IPv4:10.10.10.0/24  24324         none
                  IPv4:5.5.5.5/32     24323         none
                  IPv4:4.4.4.4/32     impl-null     24323
                  IPv4:4.4.4.4/32     impl-null     24323
                  IPv4:4.4.4.4/32     impl-null     24323
                  IPv4:4.4.4.4/32     impl-null     24323
                  IPv4:3.3.3.3/32     24322         none
                  IPv4:2.2.2.2/32     24321         none
                  IPv4:1.1.1.1/32     24320         24320
                  IPv4:1.1.1.1/32     24320         24320
                  IPv4:1.1.1.1/32     24320         24320
                  IPv4:1.1.1.1/32     24320         24320

Sent Labels :      Fec          Label          Maps To
                  IPv4:20.20.20.0/24  24325         impl-null
                  IPv4:20.20.20.0/24  24325         24322
                  IPv4:80.80.80.0/24  24328         impl-null
                  IPv4:80.80.80.0/24  24328         impl-null
                  IPv4:70.70.70.0/24  24327         impl-null
                  IPv4:70.70.70.0/24  24327         impl-null
                  IPv4:60.60.60.0/24  impl-null     none
                  IPv4:50.50.50.0/24  impl-null     none
                  IPv4:40.40.40.0/24  impl-null     none
                  IPv4:30.30.30.0/24  24326         impl-null
                  IPv4:30.30.30.0/24  24326         24322
                  IPv4:10.10.10.0/24  24324         impl-null
                  IPv4:5.5.5.5/32     impl-null     none
                  IPv4:4.4.4.4/32     24323         impl-null
                  IPv4:4.4.4.4/32     24323         24324
                  IPv4:3.3.3.3/32     24322         impl-null
                  IPv4:3.3.3.3/32     24322         24324
                  IPv4:2.2.2.2/32     24321         impl-null
                  IPv4:2.2.2.2/32     24321         24323
                  IPv4:1.1.1.1/32     24320         24320
                  IPv4:1.1.1.1/32     24320         24320
                  IPv4:1.1.1.1/32     24320         24320

Peer IP Address: 2.2.2.2
=====
Session state      : OPERATIONAL
Session role      : Active
TCP Connection     : Established
IP Address for TCP : 2.2.2.2
Interface being used : xe9
Peer LDP ID       : 2.2.2.2:0

```

```

Last Down Reason      : None
Preferred Peer LDP Password : Not Set
Adjacencies           : 40.40.40.1
                      : 2.2.2.2
Advertisement mode     : Downstream Unsolicited
Label retention mode   : Liberal
Graceful Restart      : Not Capable
Keepalive Timeout     : 30
Reconnect Interval    : 15
Configured TCP MSS    : Not configured
Applied TCP MSS       : 1460
Preferred TCP MSS     : NA
Session protection (SP) : Ready
SP duration           : infinite
Address List received  : 2.2.2.2
                      : 10.10.10.2
                      : 40.40.40.1
                      : 70.70.70.1

```

| Received Labels : | Fec | Label | Maps To |
|-------------------|--------------------|-----------|-----------|
| | IPv4:20.20.20.0/24 | 24322 | 24325 |
| | IPv4:20.20.20.0/24 | 24322 | 24325 |
| | IPv4:20.20.20.0/24 | 24322 | 24325 |
| | IPv4:20.20.20.0/24 | 24322 | 24325 |
| | IPv4:80.80.80.0/24 | 24328 | none |
| | IPv4:60.60.60.0/24 | 24327 | none |
| | IPv4:50.50.50.0/24 | 24326 | none |
| | IPv4:40.40.40.0/24 | impl-null | none |
| | IPv4:5.5.5.5/32 | 24325 | none |
| | IPv4:3.3.3.3/32 | 24324 | 24322 |
| | IPv4:3.3.3.3/32 | 24324 | 24322 |
| | IPv4:3.3.3.3/32 | 24324 | 24322 |
| | IPv4:3.3.3.3/32 | 24324 | 24322 |
| | IPv4:30.30.30.0/24 | 24323 | none |
| | IPv4:4.4.4.4/32 | 24321 | none |
| | IPv4:70.70.70.0/24 | impl-null | 24327 |
| | IPv4:70.70.70.0/24 | impl-null | 24327 |
| | IPv4:70.70.70.0/24 | impl-null | 24327 |
| | IPv4:70.70.70.0/24 | impl-null | 24327 |
| | IPv4:10.10.10.0/24 | impl-null | 24324 |
| | IPv4:10.10.10.0/24 | impl-null | 24324 |
| | IPv4:10.10.10.0/24 | impl-null | 24324 |
| | IPv4:10.10.10.0/24 | impl-null | 24324 |
| | IPv4:2.2.2.2/32 | impl-null | 24321 |
| | IPv4:2.2.2.2/32 | impl-null | 24321 |
| | IPv4:2.2.2.2/32 | impl-null | 24321 |
| | IPv4:2.2.2.2/32 | impl-null | 24321 |
| | IPv4:1.1.1.1/32 | 24320 | 24320 |
| | IPv4:1.1.1.1/32 | 24320 | 24320 |
| | IPv4:1.1.1.1/32 | 24320 | 24320 |
| | IPv4:1.1.1.1/32 | 24320 | 24320 |
| Sent Labels : | Fec | Label | Maps To |
| | IPv4:20.20.20.0/24 | 24325 | impl-null |
| | IPv4:20.20.20.0/24 | 24325 | 24322 |
| | IPv4:80.80.80.0/24 | 24328 | impl-null |
| | IPv4:80.80.80.0/24 | 24328 | impl-null |
| | IPv4:70.70.70.0/24 | 24327 | impl-null |
| | IPv4:70.70.70.0/24 | 24327 | impl-null |
| | IPv4:60.60.60.0/24 | impl-null | none |
| | IPv4:50.50.50.0/24 | impl-null | none |
| | IPv4:40.40.40.0/24 | impl-null | none |
| | IPv4:30.30.30.0/24 | 24326 | impl-null |
| | IPv4:30.30.30.0/24 | 24326 | 24322 |
| | IPv4:10.10.10.0/24 | 24324 | impl-null |
| | IPv4:5.5.5.5/32 | impl-null | none |
| | IPv4:4.4.4.4/32 | 24323 | impl-null |
| | IPv4:4.4.4.4/32 | 24323 | 24324 |
| | IPv4:3.3.3.3/32 | 24322 | impl-null |
| | IPv4:3.3.3.3/32 | 24322 | 24324 |

```

IPV4:2.2.2.2/32      24321      impl-null
IPV4:2.2.2.2/32      24321      24323
IPV4:1.1.1.1/32      24320      24320
IPV4:1.1.1.1/32      24320      24320
IPV4:1.1.1.1/32      24320      24320

PE2#show ldp session multicast count
-----
session up count    : 3
session down count  : 0
total count         : 3
-----

PE2#show ldp session targeted count
-----
session up count    : 4
session down count  : 0
total count         : 4
-----

PE2#show ldp session all count
-----
session up count    : 4
session down count  : 0
total count         : 4
-----

PE2#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity  Adjacency
-----
Multicast 3       3
Targeted  4       4
Total     7       7
-----

PE2#ping 1.1.1.1
Press CTRL+C to exit
PING 1.1.1.1 (1.1.1.1) 100(128) bytes of data.
108 bytes from 1.1.1.1: icmp_seq=1 ttl=63 time=0.445 ms
108 bytes from 1.1.1.1: icmp_seq=2 ttl=63 time=0.460 ms
108 bytes from 1.1.1.1: icmp_seq=3 ttl=63 time=0.509 ms

--- 1.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2072ms
rtt min/avg/max/mdev = 0.445/0.471/0.509/0.027 ms

```

LDP FRR with OSPF Configuration

Now that NSM, OSPF and LDP are all configured, FRR for IGP and LDP can be enabled using below configurations.

| | |
|------------------------|--|
| #configure terminal | Enter configuration mode. |
| (config)#router ospf 1 | Configure the routing process and specify the Process ID <ul style="list-style-type: none"> The Process ID should be a unique positive integer identifying the routing process. |

| | |
|--|--|
| (config-router)#fast-reroute keep-all-paths | Configure OSPF LFA-FRR to calculate the available backup path |
| (config-router)# fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp | Enables remote LFA using LDP-based tunnels as repair paths for area 0. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#fast-reroute | Configure LDP LFA-FRR to calculate the available backup path |
| (config-router)# session-protection | Enables TCP MD5-based session protection (if configured) and control-plane protection against session flaps. |
| (config-router)#exit | Exit router mode. |
| (config)#commit | Commit all the configurations |

Validation

```
PE1#show ip ospf neighbor
```

```
Total number of full neighbors: 3
```

```
OSPF process 1 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|------------|-----------|-------------|
| 2.2.2.2 | 1 | Full/DR | 00:00:31 | 10.10.10.2 | ce46/1 | 0 |
| 3.3.3.3 | 1 | Full/DR | 00:00:33 | 20.20.20.2 | ce46/4 | 0 |
| 4.4.4.4 | 1 | Full/DR | 00:00:33 | 30.30.30.2 | ce46/3 | 0 |

```
PE1#show ip ospf interface brief
```

| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
|-----------|-----|---------|---------|------|----------|-----------|--------|
| lo | 1 | 0.0.0.0 | 1 | 1 | Loopback | 0 | Up |
| ce46/1 | 1 | 0.0.0.0 | 10089 | 10 | Backup | 1 | Up |
| ce46/3 | 1 | 0.0.0.0 | 10091 | 20 | Backup | 1 | Up |
| ce46/4 | 1 | 0.0.0.0 | 10092 | 15 | Backup | 1 | Up |

```
PE1#show ip ospf database
```

```
OSPF Router with ID (1.1.1.1) (Process ID 1 VRF default)
```

```
Router Link States (Area 0.0.0.0)
```

| Link ID | ADV Router | Age | Seq# | CkSum | Link count |
|---------|------------|------|------------|--------|------------|
| 1.1.1.1 | 1.1.1.1 | 870 | 0x80000016 | 0xb291 | 4 |
| 2.2.2.2 | 2.2.2.2 | 1793 | 0x80000010 | 0x9667 | 4 |
| 3.3.3.3 | 3.3.3.3 | 890 | 0x8000001d | 0x7402 | 5 |
| 4.4.4.4 | 4.4.4.4 | 1698 | 0x80000016 | 0x357e | 4 |
| 5.5.5.5 | 5.5.5.5 | 1757 | 0x80000015 | 0x1b03 | 4 |

```
Net Link States (Area 0.0.0.0)
```

| Link ID | ADV Router | Age | Seq# | CkSum |
|------------|------------|------|------------|--------|
| 10.10.10.2 | 2.2.2.2 | 443 | 0x8000000c | 0x46b4 |
| 20.20.20.2 | 3.3.3.3 | 600 | 0x8000000b | 0xe2f2 |
| 30.30.30.2 | 4.4.4.4 | 148 | 0x8000000c | 0x7b33 |
| 40.40.40.2 | 5.5.5.5 | 487 | 0x8000000c | 0x483c |
| 50.50.50.2 | 5.5.5.5 | 347 | 0x8000000c | 0x1151 |
| 60.60.60.2 | 5.5.5.5 | 1317 | 0x8000000b | 0xdb65 |
| 70.70.70.2 | 3.3.3.3 | 1340 | 0x8000000b | 0x0635 |
| 80.80.80.1 | 4.4.4.4 | 1468 | 0x8000000b | 0xdc35 |

```

Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID      ADV Router      Age      Seq#      CkSum      Opaque ID
1.0.0.1      1.1.1.1          360      0x8000000c 0x1a02 1
1.0.0.1      2.2.2.2          73       0x8000000d 0x1cf6 1
1.0.0.1      3.3.3.3          70       0x8000000d 0x20ea 1
1.0.0.1      4.4.4.4          48       0x8000000d 0x24de 1
1.0.0.1      5.5.5.5          157      0x8000000c 0x2ad1 1
1.0.0.8      2.2.2.2          1173     0x8000000b 0x7143 8
1.0.0.10     3.3.3.3          1450     0x8000000b 0x5d50 10
1.0.0.14     3.3.3.3          1060     0x8000000b 0xf8dd 14
1.0.0.16     3.3.3.3          10       0x8000000d 0xb36b 16
1.0.0.16     5.5.5.5          1817     0x8000000c 0x977f 16
1.0.0.20     5.5.5.5          1767     0x8000000c 0x15c1 20
1.0.0.24     2.2.2.2          53       0x8000000c 0xdd7a 24
1.0.0.24     4.4.4.4          1798     0x8000000d 0x10f9 24
1.0.0.24     5.5.5.5          437      0x8000000c 0xa1a9 24
1.0.0.26     3.3.3.3          1700     0x8000000d 0x38d2 26
1.0.0.32     2.2.2.2          1743     0x8000000c 0xba49 32
1.0.0.56     4.4.4.4          1658     0x8000000c 0xab0c 56
1.0.0.62     4.4.4.4          208      0x8000000d 0x9ac9 62
1.0.0.128    1.1.1.1          1650     0x8000000d 0x4c52 128
1.0.0.132    1.1.1.1          120      0x8000000d 0xec2b 132
1.0.0.134    1.1.1.1          360      0x8000000b 0x78e0 134
PE1#
PE1#show ip route summary

-----
IP routing table name is Default-IP-Routing-Table(0)
-----
IP routing table maximum-paths : 8
Total number of IPv4 routes : 14
Total number of IPv4 paths : 14
Pending routes (due to route max reached): 0
Route Source    Networks
connected       5
ospf             9
Total           14
FIB              14

ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
Total number of IPv4 ECMP paths : 0

LFA Non ECMP statistics
-----
Total number of Routes : 9
Total number of Primary Paths : 9
Total number of Backup Paths : 9
PE1#show ip interface brief

'*' - address is assigned by dhcp client

Interface      IP-Address      Admin-Status      Link-Status
cd48            unassigned      up                 down
cd49            unassigned      up                 down
cd50            unassigned      up                 down
cd51            unassigned      up                 down
cd52            unassigned      up                 down
cd53            unassigned      up                 down
cd54            unassigned      up                 down
cd55            unassigned      up                 down
ce1             unassigned      up                 down
ce2             unassigned      up                 down
ce3             unassigned      up                 down
ce4             unassigned      up                 down
ce5             unassigned      up                 down

```



```

ce6          unassigned    up          down
ce7          unassigned    up          down
ce8          unassigned    up          down
ce9          unassigned    up          down
ce10         unassigned    up          down
ce11         unassigned    up          down
ce12         unassigned    up          down
ce13         unassigned    up          down
ce14         unassigned    up          down
ce15         unassigned    up          down
ce16         unassigned    up          down
ce17         unassigned    up          down
ce18         unassigned    up          down
ce19         unassigned    up          down
ce20         unassigned    up          down
ce21         unassigned    up          down
ce22         unassigned    up          down
ce23         unassigned    up          down
ce24         unassigned    up          down
ce25         unassigned    up          down
ce26         unassigned    up          down
ce27         unassigned    up          down
ce28         unassigned    up          down
ce29         unassigned    up          down
ce30         unassigned    up          down
ce31         unassigned    up          down
ce32         unassigned    up          down
ce33         unassigned    up          down
ce34         unassigned    up          down
ce35         unassigned    up          down
ce36         unassigned    up          down
ce37         unassigned    up          down
ce38         unassigned    up          down
ce39         unassigned    up          down
ce40         unassigned    up          down
ce41         unassigned    up          down
ce42         unassigned    up          down
ce43         unassigned    up          down
ce44         unassigned    up          up
ce45         unassigned    up          down
ce46/1       10.10.10.1      up          up
ce46/2       unassigned    up          up
ce46/3       30.30.30.1      up          up
ce46/4       20.20.20.1      up          up
ce47         unassigned    up          down
eth0         *10.16.119.123      up          up
lo           127.0.0.1      up          up
lo.management 127.0.0.1      up          up
xe0          unassigned    up          down
xe1          unassigned    up          up
xe2          unassigned    up          down
xe3          unassigned    up          down
PE1#
PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
-----
5.5.5.5      ce46/1      Passive    OPERATIONAL  30        05:29:25
2.2.2.2      ce46/1      Passive    OPERATIONAL  30        05:29:40
4.4.4.4      ce46/3      Passive    OPERATIONAL  30        05:29:40
3.3.3.3      ce46/4      Passive    OPERATIONAL  30        05:29:40
PE1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

```

E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
 ia - IS-IS inter area, E - EVPN,
 v - vrf leaked
 * - candidate default

IP Route Table for VRF "default"

```
C      1.1.1.1/32 is directly connected, lo, installed 05:30:36, last update 05:30:36 ago
O      2.2.2.2/32 [110/11] via 10.10.10.2, ce46/1, installed 05:29:42, last update 05:29:42 ago
O      3.3.3.3/32 [110/12] via 10.10.10.2, ce46/1, installed 05:29:42, last update 05:29:42 ago
O      4.4.4.4/32 [110/13] via 10.10.10.2, ce46/1, installed 05:29:22, last update 05:29:22 ago
O      5.5.5.5/32 [110/12] via 10.10.10.2, ce46/1, installed 05:29:42, last update 05:29:42 ago
C      10.10.10.0/24 is directly connected, ce46/1, installed 05:30:33, last update 05:30:33
ago
C      20.20.20.0/24 is directly connected, ce46/4, installed 05:16:01, last update 05:16:01
ago
C      30.30.30.0/24 is directly connected, ce46/3, installed 05:30:33, last update 05:30:33
ago
O      40.40.40.0/24 [110/11] via 10.10.10.2, ce46/1, installed 05:29:42, last update 05:29:42
ago
O      50.50.50.0/24 [110/12] via 10.10.10.2, ce46/1, installed 05:29:16, last update 05:29:16
ago
O      60.60.60.0/24 [110/12] via 10.10.10.2, ce46/1, installed 05:29:19, last update 05:29:19
ago
O      70.70.70.0/24 [110/11] via 10.10.10.2, ce46/1, installed 05:29:44, last update 05:29:42
ago
O      80.80.80.0/24 [110/12] via 10.10.10.2, ce46/1, installed 05:29:22, last update 05:29:22
ago
C      127.0.0.0/8 is directly connected, lo, installed 05:30:57, last update 05:30:57 ago
```

Gateway of last resort is not set

PE1#show ldp routes

| Prefix Addr | Nexthop Addr | Intf | Backup Addr | Backup |
|---------------|--------------|------------|-------------|--------------------|
| Intf | Owner | CreateTime | UpdateTime | |
| 1.1.1.1/32 | 0.0.0.0 | lo | - | - |
| | connected | 05:30:57 | - | |
| 2.2.2.2/32 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 ospf 05:30: |
| 09 05:15:31 | | | | |
| 3.3.3.3/32 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 ospf 05:30: |
| 09 05:15:31 | | | | |
| 4.4.4.4/32 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 ospf 05:30: |
| 09 05:15:31 | | | | |
| 5.5.5.5/32 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 ospf 05:30: |
| 09 05:15:31 | | | | |
| 10.10.10.0/24 | 0.0.0.0 | ce46/1 | - | - |
| | connected | 05:30:54 | - | |
| 20.20.20.0/24 | 0.0.0.0 | ce46/4 | - | - |
| | connected | 05:16:22 | - | |
| 30.30.30.0/24 | 0.0.0.0 | ce46/3 | - | - |
| | connected | 05:30:54 | - | |
| 40.40.40.0/24 | 10.10.10.2 | ce46/1 | 30.30.30.2 | ce46/3 ospf 05:30: |
| 09 05:15:31 | | | | |
| 50.50.50.0/24 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 ospf 05:30: |
| 09 05:15:31 | | | | |
| 60.60.60.0/24 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 ospf 05:30: |
| 09 05:15:31 | | | | |
| 70.70.70.0/24 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 ospf 05:30: |
| 09 05:15:31 | | | | |
| 80.80.80.0/24 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 ospf 05:30: |
| 09 05:15:31 | | | | |

PE1#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|-----|---------|----------|-----------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |

```

L> 2.2.2.2/32      1      129      -      -      -      -      -
    -              N/A      05:30:17
                        1
        Yes      3      ce46/1      No      10.10.10.2      -      -
                        128
        No      24323      ce46/4      No      20.20.20.2      -      -
L> 3.3.3.3/32      2      131      -      -      -      -      -
    -              N/A      05:30:17
                        16
        Yes      24324      ce46/1      No      10.10.10.2      -      -
                        4
        No      3      ce46/4      No      20.20.20.2      -      -
L> 4.4.4.4/32      8      135      -      -      -      -      -
    -              N/A      05:30:15
                        28
        Yes      24321      ce46/1      No      10.10.10.2      -      -
                        134
        No      24324      ce46/4      No      20.20.20.2      -      -
L> 5.5.5.5/32      3      138      -      -      -      -      -
    -              N/A      05:30:17
                        9
        Yes      24325      ce46/1      No      10.10.10.2      -      -
                        7
        No      24325      ce46/4      No      20.20.20.2      -      -
L> 40.40.40.0/24   4      141      -      -      -      -      -
    -              N/A      05:30:17
                        1
        Yes      3      ce46/1      No      10.10.10.2      -      -
                        140
        No      24326      ce46/3      No      30.30.30.2      -      -
L> 50.50.50.0/24   5      143      -      -      -      -      -
    -              N/A      05:30:17
                        19
        Yes      24326      ce46/1      No      10.10.10.2      -      -
                        4
        No      3      ce46/4      No      20.20.20.2      -      -
L> 60.60.60.0/24   9      147      -      -      -      -      -
    -              N/A      05:30:15
                        22
        Yes      24327      ce46/1      No      10.10.10.2      -      -
                        146
        No      24327      ce46/4      No      20.20.20.2      -      -
L> 70.70.70.0/24   6      150      -      -      -      -      -
    -              N/A      05:30:17
                        1
        Yes      3      ce46/1      No      10.10.10.2      -      -
                        4
        No      3      ce46/4      No      20.20.20.2      -      -
L> 80.80.80.0/24   7      152      -      -      -      -      -
    -              N/A      05:30:17
                        25
        Yes      24328      ce46/1      No      10.10.10.2      -      -
                        4
        No      3      ce46/4      No      20.20.20.2      -      -
PE1#show mpls forwarding-table 5.5.5.5/32
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code  FEC          FTN-ID  Nhlfe-ID  Tunnel-ID  Pri  Out-Label  Out-
Intf  ELC          Nexthop  Algo-Num  UpTime
L> 5.5.5.5/32      3      138      -      -      -      -      -
    -              N/A      05:30:28
                        9
        Yes      24325      ce46/1      No      10.10.10.2      -      -
                        7
        No      24325      ce46/4      No      20.20.20.2      -      -

```

```

PE1#show ldp fec
fec          fec-ipv4
PE1#show ldp fec
LSR codes    : E/N - LSR is egress/non-egress for this FEC,
               L - LSR received a label for this FEC,
               P - Primary route, B - LFA Backup route,
               R - Remote LFA Backup route,
               > - LSR will use this route for the FEC
FEC          Code    Session    Out Label    ELC    Nexthop Addr
1.1.1.1/32   NL       5.5.5.5      24320        No     connected
              NL       4.4.4.4      24320        No     connected
              NL       2.2.2.2      24320        No     connected
              NL       3.3.3.3      24320        No     connected
2.2.2.2/32   E >    non-existent  none         No     connected
              NL       5.5.5.5      24321        No     no nexthop
              NL       4.4.4.4      24321        No     no nexthop
              NLB>    3.3.3.3      24323        No     20.20.20.2
              NLP>    2.2.2.2      impl-null    No     10.10.10.2
3.3.3.3/32   NL       5.5.5.5      24322        No     no nexthop
              NL       4.4.4.4      24322        No     no nexthop
              NLP>    2.2.2.2      24324        No     10.10.10.2
              NLB>    3.3.3.3      impl-null    No     20.20.20.2
4.4.4.4/32   NL       5.5.5.5      24323        No     no nexthop
              NL       4.4.4.4      impl-null    No     no nexthop
              NLP>    2.2.2.2      24321        No     10.10.10.2
              NLB>    3.3.3.3      24324        No     20.20.20.2
5.5.5.5/32   NL       5.5.5.5      24329        No     no nexthop
              NL       4.4.4.4      24323        No     no nexthop
              NLP>    2.2.2.2      24325        No     10.10.10.2
              NLB>    3.3.3.3      24325        No     20.20.20.2
10.10.10.0/24 NL       5.5.5.5      24324        No     connected
              NL       4.4.4.4      24324        No     connected
              NL       2.2.2.2      impl-null    No     connected
              NL       3.3.3.3      24321        No     connected
              E >    non-existent  none         No     connected
20.20.20.0/24 NL       5.5.5.5      24325        No     connected
              NL       4.4.4.4      24325        No     connected
              NL       2.2.2.2      24322        No     connected
              E >    non-existent  none         No     connected
              NL       3.3.3.3      impl-null    No     connected
30.30.30.0/24 NL       5.5.5.5      24326        No     connected
              NL       4.4.4.4      impl-null    No     connected
              NL       2.2.2.2      24323        No     connected
              NL       3.3.3.3      24322        No     connected
              E >    non-existent  none         No     connected
40.40.40.0/24 NL       5.5.5.5      24330        No     no nexthop
              NLB>    4.4.4.4      24326        No     30.30.30.2
              NL       3.3.3.3      24326        No     no nexthop
              NLP>    2.2.2.2      impl-null    No     10.10.10.2
50.50.50.0/24 NL       5.5.5.5      24331        No     no nexthop
              NL       4.4.4.4      24327        No     no nexthop
              NLP>    2.2.2.2      24326        No     10.10.10.2
              NLB>    3.3.3.3      impl-null    No     20.20.20.2
60.60.60.0/24 NL       5.5.5.5      24332        No     no nexthop
              NL       4.4.4.4      impl-null    No     no nexthop
              NLP>    2.2.2.2      24327        No     10.10.10.2
              NLB>    3.3.3.3      24327        No     20.20.20.2
70.70.70.0/24 NL       5.5.5.5      24327        No     no nexthop
              NL       4.4.4.4      24328        No     no nexthop
              NLP>    2.2.2.2      impl-null    No     10.10.10.2
              NLB>    3.3.3.3      impl-null    No     20.20.20.2
80.80.80.0/24 NL       5.5.5.5      24328        No     no nexthop
              NL       4.4.4.4      impl-null    No     no nexthop
              NLP>    2.2.2.2      24328        No     10.10.10.2
              NLB>    3.3.3.3      impl-null    No     20.20.20.2

```

```
PE1#show ldp dowb
```

```
PE1#show ldp downstream
```

| FEC | Nexthop | | | | | |
|------|---------|-------|--------|------|-----------|------|
| Addr | State | Label | Req.ID | Attr | Installed | Code |

Codes: P - Primary route, B - Backup route

Session peer 5.5.5.5:

| | | | | | | |
|--|-------------|-------------|-------|------|------|----|
| 20.20.20.0/24 | connected | Established | 24325 | 0 | None | No |
| 80.80.80.0/24 | non- | | | | | |
| nh | Established | 24328 | 0 | None | No | - |
| 70.70.70.0/24 | non- | | | | | |
| nh | Established | 24327 | 0 | None | No | - |
| 60.60.60.0/24 | non- | | | | | |
| nh | Established | 24332 | 0 | None | No | - |
| 50.50.50.0/24 | non- | | | | | |
| nh | Established | 24331 | 0 | None | No | - |
| 40.40.40.0/24 | non- | | | | | |
| nh | Established | 24330 | 0 | None | No | - |
| 30.30.30.0/24 | connected | Established | 24326 | 0 | None | No |
| 10.10.10.0/24 | connected | Established | 24324 | 0 | None | No |
| 5.5.5.5/32 | non- | | | | | |
| nh | Established | 24329 | 0 | None | No | - |
| 4.4.4.4/32 | non- | | | | | |
| nh | Established | 24323 | 0 | None | No | - |
| 3.3.3.3/32 | non- | | | | | |
| nh | Established | 24322 | 0 | None | No | - |
| 2.2.2.2/32 | non- | | | | | |
| nh | Established | 24321 | 0 | None | No | - |
| 1.1.1.1/32 | connected | Established | 24320 | 0 | None | No |
| [Summary] total downstreams: 13 | | | | | | |
| [state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0 | | | | | | |
| [label] user-label: 13, impl-null: 0, expl-null: 0, invalid: 0 | | | | | | |

Codes: P - Primary route, B - Backup route

Session peer 2.2.2.2:

| | | | | | | |
|---|------------|-------------|-------|---|------|----|
| 20.20.20.0/24 | connected | Established | 24322 | 0 | None | No |
| 80.80.80.0/24 | 10.10.10.2 | Established | 24328 | 0 | None | Ye |
| s | P | | | | | |
| 60.60.60.0/24 | 10.10.10.2 | Established | 24327 | 0 | None | Ye |
| s | P | | | | | |
| 50.50.50.0/24 | 10.10.10.2 | Established | 24326 | 0 | None | Ye |
| s | P | | | | | |
| 30.30.30.0/24 | connected | Established | 24323 | 0 | None | No |
| 10.10.10.0/24 | connected | Established | impl- | | | |
| null 0 | None | No | - | | | |
| 4.4.4.4/32 | 10.10.10.2 | Established | 24321 | 0 | None | Ye |
| s | P | | | | | |
| 3.3.3.3/32 | 10.10.10.2 | Established | 24324 | 0 | None | Ye |
| s | P | | | | | |
| 1.1.1.1/32 | connected | Established | 24320 | 0 | None | No |
| 70.70.70.0/24 | 10.10.10.2 | Established | impl- | | | |
| null 0 | None | Yes | P | | | |
| 40.40.40.0/24 | 10.10.10.2 | Established | impl- | | | |
| null 0 | None | Yes | P | | | |
| 5.5.5.5/32 | 10.10.10.2 | Established | 24325 | 0 | None | Ye |
| s | P | | | | | |
| 2.2.2.2/32 | 10.10.10.2 | Established | impl- | | | |
| null 0 | None | Yes | P | | | |
| [Summary] total downstreams: 13 | | | | | | |
| [state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0 | | | | | | |
| [label] user-label: 9, impl-null: 4, expl-null: 0, invalid: 0 | | | | | | |

Codes: P - Primary route, B - Backup route

Session peer 4.4.4.4:

| | | | | | | |
|---------------|-------------|-------------|-------|------|------|----|
| 20.20.20.0/24 | connected | Established | 24325 | 0 | None | No |
| 70.70.70.0/24 | non- | | | | | |
| nh | Established | 24328 | 0 | None | No | - |
| 50.50.50.0/24 | non- | | | | | |

```

nh      Established      24327      0      None      No      -
40.40.40.0/24      30.30.30.2      Established      24326      0      None      Ye
s      B
30.30.30.0/24      connected      Established      impl-
null 0      None      No      -
10.10.10.0/24      connected      Established      24324      0      None      No
-
3.3.3.3/32      non-
nh      Established      24322      0      None      No      -
2.2.2.2/32      non-
nh      Established      24321      0      None      No      -
1.1.1.1/32      connected      Established      24320      0      None      No
-
80.80.80.0/24      non-nh      Established      impl-
null 0      None      No      -
60.60.60.0/24      non-nh      Established      impl-
null 0      None      No      -
5.5.5.5/32      non-
nh      Established      24323      0      None      No      -
4.4.4.4/32      non-nh      Established      impl-
null 0      None      No      -
[Summary] total downstreams: 13
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
[label] user-label: 9, impl-null: 4, expl-null: 0, invalid: 0

Codes: P - Primary route, B - Backup route
Session peer 3.3.3.3:
60.60.60.0/24      20.20.20.2      Established      24327      0      None      Ye
s      B
40.40.40.0/24      non-
nh      Established      24326      0      None      No      -
30.30.30.0/24      connected      Established      24322      0      None      No
-
20.20.20.0/24      connected      Established      impl-
null 0      None      No      -
10.10.10.0/24      connected      Established      24321      0      None      No
-
4.4.4.4/32      20.20.20.2      Established      24324      0      None      Ye
s      B
2.2.2.2/32      20.20.20.2      Established      24323      0      None      Ye
s      B
1.1.1.1/32      connected      Established      24320      0      None      No
-
80.80.80.0/24      20.20.20.2      Established      impl-
null 0      None      Yes      B
70.70.70.0/24      20.20.20.2      Established      impl-
null 0      None      Yes      B
50.50.50.0/24      20.20.20.2      Established      impl-
null 0      None      Yes      B
5.5.5.5/32      20.20.20.2      Established      24325      0      None      Ye
s      B
3.3.3.3/32      20.20.20.2      Established      impl-
null 0      None      Yes      B
[Summary] total downstreams: 13
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
[label] user-label: 8, impl-null: 5, expl-null: 0, invalid: 0

[Summary] total sessions: 4, total downstreams: 52

PE1#show ldp lsp
DOWNSTREAM LSP :
FEC      Nexthop
Addr      State      Label      Req.ID      Attr      Code
1.1.1.1/32      connected      Established      24320      0      None      No
-
1.1.1.1/32      connected      Established      24320      0      None      No
-
1.1.1.1/32      connected      Established      24320      0      None      No
-
1.1.1.1/32      connected      Established      24320      0      None      No

```

| | | | | | | | |
|---------------|-------------|-----------|-------------|-------|----|------|----|
| 1.1.1.1/32 | - | connected | Established | none | 0 | None | No |
| 2.2.2.2/32 | - | non- | | | | | |
| nh | Established | 24321 | 0 | None | No | - | |
| 2.2.2.2/32 | non- | | | | | | |
| nh | Established | 24321 | 0 | None | No | - | |
| 2.2.2.2/32 | 20.20.20.2 | | Established | 24323 | 0 | None | Ye |
| s | B | | | | | | |
| 2.2.2.2/32 | 10.10.10.2 | | Established | impl- | | | |
| null | 0 | None | Yes | P | | | |
| 3.3.3.3/32 | non- | | | | | | |
| nh | Established | 24322 | 0 | None | No | - | |
| 3.3.3.3/32 | non- | | | | | | |
| nh | Established | 24322 | 0 | None | No | - | |
| 3.3.3.3/32 | 10.10.10.2 | | Established | 24324 | 0 | None | Ye |
| s | P | | | | | | |
| 3.3.3.3/32 | 20.20.20.2 | | Established | impl- | | | |
| null | 0 | None | Yes | B | | | |
| 4.4.4.4/32 | non- | | | | | | |
| nh | Established | 24323 | 0 | None | No | - | |
| 4.4.4.4/32 | non-nh | | Established | impl- | | | |
| null | 0 | None | No | - | | | |
| 4.4.4.4/32 | 10.10.10.2 | | Established | 24321 | 0 | None | Ye |
| s | P | | | | | | |
| 4.4.4.4/32 | 20.20.20.2 | | Established | 24324 | 0 | None | Ye |
| s | B | | | | | | |
| 5.5.5.5/32 | non- | | | | | | |
| nh | Established | 24329 | 0 | None | No | - | |
| 5.5.5.5/32 | non- | | | | | | |
| nh | Established | 24323 | 0 | None | No | - | |
| 5.5.5.5/32 | 10.10.10.2 | | Established | 24325 | 0 | None | Ye |
| s | P | | | | | | |
| 5.5.5.5/32 | 20.20.20.2 | | Established | 24325 | 0 | None | Ye |
| s | B | | | | | | |
| 10.10.10.0/24 | connected | | Established | 24324 | 0 | None | No |
| 10.10.10.0/24 | connected | | Established | 24324 | 0 | None | No |
| 10.10.10.0/24 | connected | | Established | impl- | | | |
| null | 0 | None | No | - | | | |
| 10.10.10.0/24 | connected | | Established | 24321 | 0 | None | No |
| 10.10.10.0/24 | connected | | Established | none | 0 | None | No |
| 20.20.20.0/24 | connected | | Established | 24325 | 0 | None | No |
| 20.20.20.0/24 | connected | | Established | 24325 | 0 | None | No |
| 20.20.20.0/24 | connected | | Established | 24322 | 0 | None | No |
| 20.20.20.0/24 | connected | | Established | none | 0 | None | No |
| 20.20.20.0/24 | connected | | Established | impl- | | | |
| null | 0 | None | No | - | | | |
| 30.30.30.0/24 | connected | | Established | 24326 | 0 | None | No |
| 30.30.30.0/24 | connected | | Established | impl- | | | |
| null | 0 | None | No | - | | | |
| 30.30.30.0/24 | connected | | Established | 24323 | 0 | None | No |
| 30.30.30.0/24 | connected | | Established | 24322 | 0 | None | No |
| 30.30.30.0/24 | connected | | Established | none | 0 | None | No |
| 40.40.40.0/24 | non- | | | | | | |
| nh | Established | 24330 | 0 | None | No | - | |
| 40.40.40.0/24 | 30.30.30.2 | | Established | 24326 | 0 | None | Ye |
| s | B | | | | | | |
| 40.40.40.0/24 | non- | | | | | | |
| nh | Established | 24326 | 0 | None | No | - | |

| | | | | | | | | |
|----------------|-------------|-------------|--------|------|---------|------|---|----|
| 40.40.40.0/24 | 10.10.10.2 | Established | impl- | | | | | |
| null 0 | None | Yes | P | | | | | |
| 50.50.50.0/24 | non- | | | | | | | |
| nh | Established | 24331 | 0 | None | No | - | | |
| 50.50.50.0/24 | non- | | | | | | | |
| nh | Established | 24327 | 0 | None | No | - | | |
| 50.50.50.0/24 | 10.10.10.2 | Established | 24326 | 0 | None | | | Ye |
| s | P | | | | | | | |
| 50.50.50.0/24 | 20.20.20.2 | Established | impl- | | | | | |
| null 0 | None | Yes | B | | | | | |
| 60.60.60.0/24 | non- | | | | | | | |
| nh | Established | 24332 | 0 | None | No | - | | |
| 60.60.60.0/24 | non-nh | Established | impl- | | | | | |
| null 0 | None | No | - | | | | | |
| 60.60.60.0/24 | 10.10.10.2 | Established | 24327 | 0 | None | | | Ye |
| s | P | | | | | | | |
| 60.60.60.0/24 | 20.20.20.2 | Established | 24327 | 0 | None | | | Ye |
| s | B | | | | | | | |
| 70.70.70.0/24 | non- | | | | | | | |
| nh | Established | 24327 | 0 | None | No | - | | |
| 70.70.70.0/24 | non- | | | | | | | |
| nh | Established | 24328 | 0 | None | No | - | | |
| 70.70.70.0/24 | 10.10.10.2 | Established | impl- | | | | | |
| null 0 | None | Yes | P | | | | | |
| 70.70.70.0/24 | 20.20.20.2 | Established | impl- | | | | | |
| null 0 | None | Yes | B | | | | | |
| 80.80.80.0/24 | non- | | | | | | | |
| nh | Established | 24328 | 0 | None | No | - | | |
| 80.80.80.0/24 | non-nh | Established | impl- | | | | | |
| null 0 | None | No | - | | | | | |
| 80.80.80.0/24 | 10.10.10.2 | Established | 24328 | 0 | None | | | Ye |
| s | P | | | | | | | |
| 80.80.80.0/24 | 20.20.20.2 | Established | impl- | | | | | |
| null 0 | None | Yes | B | | | | | |
| UPSTREAM LSP : | | | | | | | | |
| FEC | State | Label | Req.ID | Attr | | | | |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes | 1 | |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes | 1 | |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes | 1 | |
| 2.2.2.2/32 | Established | 24320 | 0 | None | skipped | /no | 3 | |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes | 3 | |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes | 3 | |
| 2.2.2.2/32 | Established | 24320 | 0 | None | skipped | /no | 3 | |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes | 3 | |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes | 3 | |
| 3.3.3.3/32 | Established | 24321 | 0 | None | skipped | /no | 3 | |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes | 3 | |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes | 3 | |
| 3.3.3.3/32 | Established | 24321 | 0 | None | skipped | /no | 3 | |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes | 3 | |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes | 3 | |
| 4.4.4.4/32 | Established | 24322 | 0 | None | skipped | /no | 3 | |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes | 3 | |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes | 3 | |
| 4.4.4.4/32 | Established | 24322 | 0 | None | skipped | /no | 3 | |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes | 3 | |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes | 3 | |
| 5.5.5.5/32 | Established | 24323 | 0 | None | skipped | /no | 3 | |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes | 3 | |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes | 3 | |
| 5.5.5.5/32 | Established | 24323 | 0 | None | skipped | /no | 3 | |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes | 3 | |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes | 3 | |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes | 1 | |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes | 1 | |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes | 1 | |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes | 1 | |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes | 1 | |

| | | | | | | | |
|---------------|-------------|-----------|---|------|---------|------|---|
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | skipped | /no | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | skipped | /no | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | skipped | /no | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | skipped | /no | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | skipped | /no | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | skipped | /no | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |

```
PE1#show ldp fec
```

```
fec          fec-ipv4
```

```
PE1#show ldp fec prefix 5.5.5.5/32
```

```
LSR codes    : E/N - LSR is egress/non-egress for this FEC,
               L - LSR received a label for this FEC,
               P - Primary route, B - LFA Backup route,
               R - Remote LFA Backup route,
               > - LSR will use this route for the FEC
```

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|------------|------|---------|-----------|-----|--------------|
| 5.5.5.5/32 | NL | 5.5.5.5 | 24329 | No | no nexthop |
| | NL | 4.4.4.4 | 24323 | No | no nexthop |
| | NLP> | 2.2.2.2 | 24325 | No | 10.10.10.2 |
| | NLB> | 3.3.3.3 | 24325 | No | 20.20.20.2 |

```
PE1#show mpls ftn-table
```

```
Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed
```

```
CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27
```

```
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 refcount: 1
```

```
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 7, out intf: ce46/1, out label: 3
Nexthop addr: 10.10.10.2      cross connect ix: 3, op code: Push
```

```
Backup Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 128 bypass ftn-ix: 0
```

```
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 128, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/4, out label: 24323
Nexthop addr: 20.20.20.2      cross connect ix: 1, op code: Push
```

```
Primary FTN entry with FEC: 3.3.3.3/32, id: 2, row status: Active, Tunnel-Policy: N/A, State: Installed
```

```
CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27
```

```
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
```

```

Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 16 refcount: 1
    Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 16, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24324
  Nexthop addr: 10.10.10.2          cross connect ix: 4, op code: Push

  Backup Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3
  Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 05:33:02, UpTime: 05:33:02, LastUpdate: 05:18:27
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 28 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 28, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24321
    Nexthop addr: 10.10.10.2          cross connect ix: 9, op code: Push

    Backup Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 134 bypass ftn-ix: 0
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 134, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/4, out label:
24324
    Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 9 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label: 24325
    Nexthop addr: 10.10.10.2          cross connect ix: 4, op code: Push

    Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 bypass ftn-ix: 0
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: ce46/4, out label: 24325
    Nexthop addr: 20.20.20.2          cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 40.40.40.0/24, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 refcount: 1
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 7, out intf: ce46/1, out label: 3
    Nexthop addr: 10.10.10.2          cross connect ix: 3, op code: Push

    Backup Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 140 bypass ftn-ix: 0
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 140, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/3, out label:
24326
    Nexthop addr: 30.30.30.2          cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 50.50.50.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed

```

```
CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 19 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 19, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24326
Nexthop addr: 10.10.10.2          cross connect ix: 6, op code: Push

Backup Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 60.60.60.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 05:33:02, UpTime: 05:33:02, LastUpdate: 05:18:27
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 22 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 22, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24327
Nexthop addr: 10.10.10.2          cross connect ix: 7, op code: Push

Backup Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 146 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 146, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/4, out label:
24327
Nexthop addr: 20.20.20.2          cross connect ix: 9, op code: Push

Primary FTN entry with FEC: 70.70.70.0/24, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 7, out intf: ce46/1, out label: 3
Nexthop addr: 10.10.10.2          cross connect ix: 3, op code: Push

Backup Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 80.80.80.0/24, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 05:33:04, UpTime: 05:33:04, LastUpdate: 05:18:27
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 25 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 25, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24328
Nexthop addr: 10.10.10.2          cross connect ix: 8, op code: Push

Backup Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push
```

```

PE1#show mpls ftn-table 5.5.5.5/32
Primary FTN entry with FEC: 5.5.5.5/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 05:33:57, UpTime: 05:33:57, LastUpdate: 05:19:20
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 9 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label: 24325
      Nexthop addr: 10.10.10.2          cross connect ix: 4, op code: Push

    Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 bypass ftn-ix: 0
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: ce46/4, out label: 24325
      Nexthop addr: 20.20.20.2          cross connect ix: 4, op code: Push

PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
      K - CLI ILM, T - MPLS-TP, s - Stitched ILM
      S - SNMP, L - LDP, R - RSVP, C - CRLDP
      B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
      O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
      P - SR Policy,          U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled
Code  FEC/VRF/L2CKT  ILM-ID  In-Label  Out-Label  In-Intf  Out-
Intf/VRF  Nexthop  pri  Algo-Num  UpTime  UPStr  peers
L>  4.4.4.4/32      2      24322     24321     N/A      ce46/1      10.10.10.2
      Yes N/A      05:34:09  2
L>  2.2.2.2/32      1      24320     3         N/A      ce46/1      10.10.10.2
      Yes N/A      05:34:09  2
L>  3.3.3.3/32      5      24321     24324     N/A      ce46/1      10.10.10.2
      Yes N/A      05:34:09  2
L>  50.50.50.0/24   6      24325     24326     N/A      ce46/1      10.10.10.2
      Yes N/A      05:34:09  2
L>  40.40.40.0/24   3      24324     3         N/A      ce46/1      10.10.10.2
      Yes N/A      05:34:09  2
      24324     24326     N/A      ce46/3      30.30.30.2
      No -      -
L>  5.5.5.5/32      9      24323     24325     N/A      ce46/1      10.10.10.2
      Yes N/A      05:34:06  2
      24323     24325     N/A      ce46/4      20.20.20.2
      No -      -
L>  70.70.70.0/24   8      24327     3         N/A      ce46/1      10.10.10.2
      Yes N/A      05:34:08  2
      24327     3         N/A      ce46/4      20.20.20.2
      No -      -
L>  60.60.60.0/24   4      24326     24327     N/A      ce46/1      10.10.10.2
      Yes N/A      05:34:09  2
L>  80.80.80.0/24   7      24328     24328     N/A      ce46/1      10.10.10.2
      Yes N/A      05:34:05  2

PE1#show ldp igp sync
Interface      Link-State  LDP-Enabled  IGP-Sync  Peer-IP      Session-State  Sync-Delay-
Time/Remaining-Delay-Time
ce46/1         Up          Yes          Enabled   2.2.2.2      Achieved       Not-Configured / Not-
Running
ce46/3         Up          Yes          Enabled   4.4.4.4      Achieved       Not-Configured / Not-
Running
ce46/4         Up          Yes          Enabled   3.3.3.3      Achieved       Not-Configured / Not-
Running
PE1#show ldp session multicast count
-----
session up count   : 3
session down count : 0
total count        : 3
-----
PE1#show ldp session targeted count

```

```

-----
session up count   : 4
session down count : 0
total count        : 4
-----

PE1#show ldp session all count
-----
session up count   : 4
session down count : 0
total count        : 4
-----

PE1#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity  Adjacency
-----
Multicast 3       3
Targeted  4       4
Total     7       7
-----

PE1#
Vty connection is timed out.
root@PE1:~#
root@PE1:~# cmlsh

OcNOS version UFI_S9600-56DX-OcNOS-6.6.1.140-SP_PLUS_Q2-MR 07/01/2025 14:47:46
PE1>
PE1>
PE1>en
PE1#
PE1#
PE1#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area ,p - stale info, E -
EVPN
       * - candidate default

IP Route Table for VRF "default"
O      2.2.2.2/32 [110/11] via 10.10.10.2, ce46/1, installed 06:03:44, last update 06:03:44 ago
        [FRR-NH] via 20.20.20.2, ce46/4

O      3.3.3.3/32 [110/12] via 10.10.10.2, ce46/1, installed 06:03:44, last update 06:03:44 ago
        [FRR-NH] via 20.20.20.2, ce46/4

O      4.4.4.4/32 [110/13] via 10.10.10.2, ce46/1, installed 06:03:24, last update 06:03:24 ago
        [FRR-NH] via 20.20.20.2, ce46/4

O      5.5.5.5/32 [110/12] via 10.10.10.2, ce46/1, installed 06:03:44, last update 06:03:44 ago
        [FRR-NH] via 20.20.20.2, ce46/4

O      40.40.40.0/24 [110/11] via 10.10.10.2, ce46/1, installed 06:03:44, last update 06:03:44 ago
        [FRR-NH] via 30.30.30.2, ce46/3

O      50.50.50.0/24 [110/12] via 10.10.10.2, ce46/1, installed 06:03:18, last update 06:03:18 ago
        [FRR-NH] via 20.20.20.2, ce46/4

O      60.60.60.0/24 [110/12] via 10.10.10.2, ce46/1, installed 06:03:21, last update 06:03:21 ago
        [FRR-NH] via 20.20.20.2, ce46/4

O      70.70.70.0/24 [110/11] via 10.10.10.2, ce46/1, installed 06:03:46, last update 06:03:44 ago
        [FRR-NH] via 20.20.20.2, ce46/4

O      80.80.80.0/24 [110/12] via 10.10.10.2, ce46/1, installed 06:03:24, last update 06:03:24 ago
        [FRR-NH] via 20.20.20.2, ce46/4

```

```
PE1#show ip ospf route fast-reroute
```

```
OSPF process 1:
```

```
Codes: C - connected, D - Discard, O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
OSPF LFA attributes:
```

```
P - Primary, SP - Secondary-Path, LP - Link Protecting,
```

```
NP - Node Protecting, BID - Broadcast Link Protecting
```

```
DP - Downstream Protecting
```

```
O 2.2.2.2/32 [11] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 20.20.20.2, ce46/4, Area 0.0.0.0
  Attributes       :
  Metric           : [17] ,SP ,BID ,DP
O 3.3.3.3/32 [12] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 20.20.20.2, ce46/4, Area 0.0.0.0
  Attributes       :
  Metric           : [16] ,SP ,NP ,BID ,DP
O 4.4.4.4/32 [13] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 20.20.20.2, ce46/4, Area 0.0.0.0
  Attributes       :
  Metric           : [17] ,SP ,NP ,BID ,DP
O 5.5.5.5/32 [12] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 20.20.20.2, ce46/4, Area 0.0.0.0
  Attributes       :
  Metric           : [17] ,SP ,NP ,BID ,DP
O 20.20.20.0/24 [12] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 20.20.20.2, ce46/4, Area 0.0.0.0
  Attributes       :
  Metric           : [16] ,SP ,NP ,BID ,DP
O 30.30.30.0/24 [13] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 30.30.30.2, ce46/3, Area 0.0.0.0
  Attributes       :
  Metric           : [21] ,SP ,NP ,BID ,DP
O 40.40.40.0/24 [11] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 30.30.30.2, ce46/3, Area 0.0.0.0
  Attributes       :
  Metric           : [22] ,SP ,NP ,BID ,DP
O 50.50.50.0/24 [12] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 20.20.20.2, ce46/4, Area 0.0.0.0
  Attributes       :
  Metric           : [16] ,SP ,NP ,BID ,DP
O 60.60.60.0/24 [12] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 20.20.20.2, ce46/4, Area 0.0.0.0
  Attributes       :
  Metric           : [17] ,SP ,NP ,BID ,DP
O 70.70.70.0/24 [11] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 20.20.20.2, ce46/4, Area 0.0.0.0
  Attributes       :
  Metric           : [16] ,SP ,NP ,BID ,DP
O 80.80.80.0/24 [12] via 10.10.10.2, ce46/1, Area 0.0.0.0
  Backup path      : via 20.20.20.2, ce46/4, Area 0.0.0.0
  Attributes       :
  Metric           : [16] ,SP ,NP ,BID ,DP
```

```
PE1#show ldp routes
```

| Prefix Addr | NextHop Addr | Intf | Backup Addr | Backup | | |
|-------------|--------------|------------|-------------|--------|------|--------|
| Intf | Owner | CreateTime | UpdateTime | | | |
| 1.1.1.1/32 | 0.0.0.0 | lo | - | - | | |
| | connected | 06:05:12 | - | | | |
| 2.2.2.2/32 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 | ospf | 06:04: |
| 24 05:49:46 | | | | | | |
| 3.3.3.3/32 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 | ospf | 06:04: |
| 24 05:49:46 | | | | | | |
| 4.4.4.4/32 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 | ospf | 06:04: |
| 24 05:49:46 | | | | | | |
| 5.5.5.5/32 | 10.10.10.2 | ce46/1 | 20.20.20.2 | ce46/4 | ospf | 06:04: |
| 24 05:49:46 | | | | | | |

```

10.10.10.0/24      0.0.0.0      ce46/1      -      -
connected 06:05:09 -
20.20.20.0/24      0.0.0.0      ce46/4      -      -
connected 05:50:37 -
30.30.30.0/24      0.0.0.0      ce46/3      -      -
connected 06:05:09 -
40.40.40.0/24      10.10.10.2   ce46/1      30.30.30.2  ce46/3      ospf      06:04:
24 05:49:46
50.50.50.0/24      10.10.10.2   ce46/1      20.20.20.2  ce46/4      ospf      06:04:
24 05:49:46
60.60.60.0/24      10.10.10.2   ce46/1      20.20.20.2  ce46/4      ospf      06:04:
24 05:49:46
70.70.70.0/24      10.10.10.2   ce46/1      20.20.20.2  ce46/4      ospf      06:04:
24 05:49:46
80.80.80.0/24      10.10.10.2   ce46/1      20.20.20.2  ce46/4      ospf      06:04:
24 05:49:46
PE1#show ldp fec
LSR codes : E/N - LSR is egress/non-egress for this FEC,
            L - LSR received a label for this FEC,
            P - Primary route, B - LFA Backup route,
            R - Remote LFA Backup route,
            > - LSR will use this route for the FEC
FEC          Code   Session      Out Label   ELC   Nexthop Addr
1.1.1.1/32   NL    5.5.5.5     24320      No    connected
              NL    4.4.4.4     24320      No    connected
              NL    2.2.2.2     24320      No    connected
              NL    3.3.3.3     24320      No    connected
              E >   non-existent none       No    connected
2.2.2.2/32   NL    5.5.5.5     24321      No    no nexthop
              NL    4.4.4.4     24321      No    no nexthop
              NLB>  3.3.3.3     24323      No    20.20.20.2
              NLP>  2.2.2.2     impl-null  No    10.10.10.2
3.3.3.3/32   NL    5.5.5.5     24322      No    no nexthop
              NL    4.4.4.4     24322      No    no nexthop
              NLP>  2.2.2.2     24324      No    10.10.10.2
              NLB>  3.3.3.3     impl-null  No    20.20.20.2
4.4.4.4/32   NL    5.5.5.5     24323      No    no nexthop
              NL    4.4.4.4     impl-null  No    no nexthop
              NLP>  2.2.2.2     24321      No    10.10.10.2
              NLB>  3.3.3.3     24324      No    20.20.20.2
5.5.5.5/32   NL    5.5.5.5     24329      No    no nexthop
              NL    4.4.4.4     24323      No    no nexthop
              NLP>  2.2.2.2     24325      No    10.10.10.2
              NLB>  3.3.3.3     24325      No    20.20.20.2
10.10.10.0/24 NL    5.5.5.5     24324      No    connected
              NL    4.4.4.4     24324      No    connected
              NL    2.2.2.2     impl-null  No    connected
              NL    3.3.3.3     24321      No    connected
              E >   non-existent none       No    connected
20.20.20.0/24 NL    5.5.5.5     24325      No    connected
              NL    4.4.4.4     24325      No    connected
              NL    2.2.2.2     24322      No    connected
              E >   non-existent none       No    connected
              NL    3.3.3.3     impl-null  No    connected
30.30.30.0/24 NL    5.5.5.5     24326      No    connected
              NL    4.4.4.4     impl-null  No    connected
              NL    2.2.2.2     24323      No    connected
              NL    3.3.3.3     24322      No    connected
              E >   non-existent none       No    connected
40.40.40.0/24 NL    5.5.5.5     24330      No    no nexthop
              NLB>  4.4.4.4     24326      No    30.30.30.2
              NL    3.3.3.3     24326      No    no nexthop
              NLP>  2.2.2.2     impl-null  No    10.10.10.2
50.50.50.0/24 NL    5.5.5.5     24331      No    no nexthop
              NL    4.4.4.4     24327      No    no nexthop
              NLP>  2.2.2.2     24326      No    10.10.10.2
              NLB>  3.3.3.3     impl-null  No    20.20.20.2
60.60.60.0/24 NL    5.5.5.5     24332      No    no nexthop
              NL    4.4.4.4     impl-null  No    no nexthop

```

```

NLP> 2.2.2.2 24327 No 10.10.10.2
NLB> 3.3.3.3 24327 No 20.20.20.2
70.70.70.0/24 NL 5.5.5.5 24327 No no nexthop
NL 4.4.4.4 24328 No no nexthop
NLP> 2.2.2.2 impl-null No 10.10.10.2
NLB> 3.3.3.3 impl-null No 20.20.20.2
80.80.80.0/24 NL 5.5.5.5 24328 No no nexthop
NL 4.4.4.4 impl-null No no nexthop
NLP> 2.2.2.2 24328 No 10.10.10.2
NLB> 3.3.3.3 impl-null No 20.20.20.2

PE1#show ldp dowb
PE1#show ldp downstream
FEC Nexthop
Addr State Label Req.ID Attr Installed Code
Codes: P - Primary route, B - Backup route
Session peer 5.5.5.5:
20.20.20.0/24 connected Established 24325 0 None No
-
80.80.80.0/24 non-
nh Established 24328 0 None No -
70.70.70.0/24 non-
nh Established 24327 0 None No -
60.60.60.0/24 non-
nh Established 24332 0 None No -
50.50.50.0/24 non-
nh Established 24331 0 None No -
40.40.40.0/24 non-
nh Established 24330 0 None No -
30.30.30.0/24 connected Established 24326 0 None No
-
10.10.10.0/24 connected Established 24324 0 None No
-
5.5.5.5/32 non-
nh Established 24329 0 None No -
4.4.4.4/32 non-
nh Established 24323 0 None No -
3.3.3.3/32 non-
nh Established 24322 0 None No -
2.2.2.2/32 non-
nh Established 24321 0 None No -
1.1.1.1/32 connected Established 24320 0 None No
-
[Summary] total downstreams: 13
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
[label] user-label: 13, impl-null: 0, expl-null: 0, invalid: 0

Codes: P - Primary route, B - Backup route
Session peer 2.2.2.2:
20.20.20.0/24 connected Established 24322 0 None No
-
80.80.80.0/24 10.10.10.2 Established 24328 0 None Ye
s P
60.60.60.0/24 10.10.10.2 Established 24327 0 None Ye
s P
50.50.50.0/24 10.10.10.2 Established 24326 0 None Ye
s P
30.30.30.0/24 connected Established 24323 0 None No
-
10.10.10.0/24 connected Established impl-
null 0 None No -
4.4.4.4/32 10.10.10.2 Established 24321 0 None Ye
s P
3.3.3.3/32 10.10.10.2 Established 24324 0 None Ye
s P
1.1.1.1/32 connected Established 24320 0 None No
-
70.70.70.0/24 10.10.10.2 Established impl-
null 0 None Yes P
40.40.40.0/24 10.10.10.2 Established impl-
null 0 None Yes P

```


| | | | | | | | |
|---|---------------|------------|-------------|-------|----|------|----|
| s | 5.5.5.5/32 | 10.10.10.2 | Established | 24325 | 0 | None | Ye |
| P | 2.2.2.2/32 | 10.10.10.2 | Established | impl- | | | |
| nh | 0 | None | Yes | P | | | |
| [Summary] total downstreams: 13 | | | | | | | |
| [state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0 | | | | | | | |
| [label] user-label: 9, impl-null: 4, expl-null: 0, invalid: 0 | | | | | | | |
| Codes: P - Primary route, B - Backup route | | | | | | | |
| Session peer 4.4.4.4: | | | | | | | |
| | 20.20.20.0/24 | connected | Established | 24325 | 0 | None | No |
| - | 70.70.70.0/24 | non- | | | | | |
| nh | Established | 24328 | 0 | None | No | - | |
| nh | 50.50.50.0/24 | non- | | | | | |
| nh | Established | 24327 | 0 | None | No | - | |
| s | 40.40.40.0/24 | 30.30.30.2 | Established | 24326 | 0 | None | Ye |
| B | 30.30.30.0/24 | connected | Established | impl- | | | |
| nh | 0 | None | No | - | | | |
| | 10.10.10.0/24 | connected | Established | 24324 | 0 | None | No |
| - | 3.3.3.3/32 | non- | | | | | |
| nh | Established | 24322 | 0 | None | No | - | |
| nh | 2.2.2.2/32 | non- | | | | | |
| nh | Established | 24321 | 0 | None | No | - | |
| | 1.1.1.1/32 | connected | Established | 24320 | 0 | None | No |
| - | 80.80.80.0/24 | non-nh | Established | impl- | | | |
| nh | 0 | None | No | - | | | |
| | 60.60.60.0/24 | non-nh | Established | impl- | | | |
| nh | 0 | None | No | - | | | |
| | 5.5.5.5/32 | non- | | | | | |
| nh | Established | 24323 | 0 | None | No | - | |
| | 4.4.4.4/32 | non-nh | Established | impl- | | | |
| nh | 0 | None | No | - | | | |
| [Summary] total downstreams: 13 | | | | | | | |
| [state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0 | | | | | | | |
| [label] user-label: 9, impl-null: 4, expl-null: 0, invalid: 0 | | | | | | | |
| Codes: P - Primary route, B - Backup route | | | | | | | |
| Session peer 3.3.3.3: | | | | | | | |
| | 60.60.60.0/24 | 20.20.20.2 | Established | 24327 | 0 | None | Ye |
| B | 40.40.40.0/24 | non- | | | | | |
| nh | Established | 24326 | 0 | None | No | - | |
| | 30.30.30.0/24 | connected | Established | 24322 | 0 | None | No |
| - | 20.20.20.0/24 | connected | Established | impl- | | | |
| nh | 0 | None | No | - | | | |
| | 10.10.10.0/24 | connected | Established | 24321 | 0 | None | No |
| - | 4.4.4.4/32 | 20.20.20.2 | Established | 24324 | 0 | None | Ye |
| B | 2.2.2.2/32 | 20.20.20.2 | Established | 24323 | 0 | None | Ye |
| s | 1.1.1.1/32 | connected | Established | 24320 | 0 | None | No |
| - | 80.80.80.0/24 | 20.20.20.2 | Established | impl- | | | |
| nh | 0 | None | Yes | B | | | |
| | 70.70.70.0/24 | 20.20.20.2 | Established | impl- | | | |
| nh | 0 | None | Yes | B | | | |
| | 50.50.50.0/24 | 20.20.20.2 | Established | impl- | | | |
| nh | 0 | None | Yes | B | | | |
| | 5.5.5.5/32 | 20.20.20.2 | Established | 24325 | 0 | None | Ye |
| B | 3.3.3.3/32 | 20.20.20.2 | Established | impl- | | | |
| nh | 0 | None | Yes | B | | | |
| [Summary] total downstreams: 13 | | | | | | | |
| [state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0 | | | | | | | |

```
[label] user-label: 8, impl-null: 5, expl-null: 0, invalid: 0
```

```
[Summary] total sessions: 4, total downstreams: 52
```

```
PE1#show ldp lsp
```

```
DOWNSTREAM LSP :
```

| FEC | State | Nexthop Label | Req.ID | Attr | Code | | | |
|---------------|-------------|---------------|-------------|-------|------|------|--|----|
| Addr | | | | | | | | |
| 1.1.1.1/32 | | connected | Established | 24320 | 0 | None | | No |
| - | | | | | | | | |
| 1.1.1.1/32 | | connected | Established | 24320 | 0 | None | | No |
| - | | | | | | | | |
| 1.1.1.1/32 | | connected | Established | 24320 | 0 | None | | No |
| - | | | | | | | | |
| 1.1.1.1/32 | | connected | Established | 24320 | 0 | None | | No |
| - | | | | | | | | |
| 1.1.1.1/32 | | connected | Established | none | 0 | None | | No |
| - | | | | | | | | |
| 2.2.2.2/32 | | non- | | | | | | |
| nh | Established | 24321 | 0 | None | No | - | | |
| 2.2.2.2/32 | | non- | | | | | | |
| nh | Established | 24321 | 0 | None | No | - | | |
| 2.2.2.2/32 | | 20.20.20.2 | Established | 24323 | 0 | None | | Ye |
| s | B | | | | | | | |
| 2.2.2.2/32 | | 10.10.10.2 | Established | impl- | | | | |
| null 0 | None | | Yes | P | | | | |
| 3.3.3.3/32 | | non- | | | | | | |
| nh | Established | 24322 | 0 | None | No | - | | |
| 3.3.3.3/32 | | non- | | | | | | |
| nh | Established | 24322 | 0 | None | No | - | | |
| 3.3.3.3/32 | | 10.10.10.2 | Established | 24324 | 0 | None | | Ye |
| s | P | | | | | | | |
| 3.3.3.3/32 | | 20.20.20.2 | Established | impl- | | | | |
| null 0 | None | | Yes | B | | | | |
| 4.4.4.4/32 | | non- | | | | | | |
| nh | Established | 24323 | 0 | None | No | - | | |
| 4.4.4.4/32 | | non-nh | Established | impl- | | | | |
| null 0 | None | | No | - | | | | |
| 4.4.4.4/32 | | 10.10.10.2 | Established | 24321 | 0 | None | | Ye |
| s | P | | | | | | | |
| 4.4.4.4/32 | | 20.20.20.2 | Established | 24324 | 0 | None | | Ye |
| s | B | | | | | | | |
| 5.5.5.5/32 | | non- | | | | | | |
| nh | Established | 24329 | 0 | None | No | - | | |
| 5.5.5.5/32 | | non- | | | | | | |
| nh | Established | 24323 | 0 | None | No | - | | |
| 5.5.5.5/32 | | 10.10.10.2 | Established | 24325 | 0 | None | | Ye |
| s | P | | | | | | | |
| 5.5.5.5/32 | | 20.20.20.2 | Established | 24325 | 0 | None | | Ye |
| s | B | | | | | | | |
| 10.10.10.0/24 | | connected | Established | 24324 | 0 | None | | No |
| - | | | | | | | | |
| 10.10.10.0/24 | | connected | Established | 24324 | 0 | None | | No |
| - | | | | | | | | |
| 10.10.10.0/24 | | connected | Established | impl- | | | | |
| null 0 | None | | No | - | | | | |
| 10.10.10.0/24 | | connected | Established | 24321 | 0 | None | | No |
| - | | | | | | | | |
| 10.10.10.0/24 | | connected | Established | none | 0 | None | | No |
| - | | | | | | | | |
| 20.20.20.0/24 | | connected | Established | 24325 | 0 | None | | No |
| - | | | | | | | | |
| 20.20.20.0/24 | | connected | Established | 24325 | 0 | None | | No |
| - | | | | | | | | |
| 20.20.20.0/24 | | connected | Established | 24322 | 0 | None | | No |
| - | | | | | | | | |
| 20.20.20.0/24 | | connected | Established | none | 0 | None | | No |
| - | | | | | | | | |
| 20.20.20.0/24 | | connected | Established | impl- | | | | |
| null 0 | None | | No | - | | | | |

| | | | | | | |
|----------------|-------------|-------------|--------|------|---------|--------|
| 30.30.30.0/24 | connected | Established | 24326 | 0 | None | No |
| 30.30.30.0/24 | connected | Established | impl- | | | |
| 30.30.30.0/24 | connected | Established | 24323 | 0 | None | No |
| 30.30.30.0/24 | connected | Established | 24322 | 0 | None | No |
| 30.30.30.0/24 | connected | Established | none | 0 | None | No |
| 40.40.40.0/24 | non- | | | | | |
| nh | Established | 24330 | 0 | None | No | - |
| 40.40.40.0/24 | 30.30.30.2 | Established | 24326 | 0 | None | Ye |
| 40.40.40.0/24 | non- | | | | | |
| nh | Established | 24326 | 0 | None | No | - |
| 40.40.40.0/24 | 10.10.10.2 | Established | impl- | | | |
| 30.30.30.0/24 | None | Yes | P | | | |
| 50.50.50.0/24 | non- | | | | | |
| nh | Established | 24331 | 0 | None | No | - |
| 50.50.50.0/24 | non- | | | | | |
| nh | Established | 24327 | 0 | None | No | - |
| 50.50.50.0/24 | 10.10.10.2 | Established | 24326 | 0 | None | Ye |
| 50.50.50.0/24 | 20.20.20.2 | Established | impl- | | | |
| 30.30.30.0/24 | None | Yes | B | | | |
| 60.60.60.0/24 | non- | | | | | |
| nh | Established | 24332 | 0 | None | No | - |
| 60.60.60.0/24 | non-nh | Established | impl- | | | |
| 30.30.30.0/24 | None | No | - | | | |
| 60.60.60.0/24 | 10.10.10.2 | Established | 24327 | 0 | None | Ye |
| 60.60.60.0/24 | 20.20.20.2 | Established | 24327 | 0 | None | Ye |
| 70.70.70.0/24 | non- | | | | | |
| nh | Established | 24327 | 0 | None | No | - |
| 70.70.70.0/24 | non- | | | | | |
| nh | Established | 24328 | 0 | None | No | - |
| 70.70.70.0/24 | 10.10.10.2 | Established | impl- | | | |
| 30.30.30.0/24 | None | Yes | P | | | |
| 70.70.70.0/24 | 20.20.20.2 | Established | impl- | | | |
| 30.30.30.0/24 | None | Yes | B | | | |
| 80.80.80.0/24 | non- | | | | | |
| nh | Established | 24328 | 0 | None | No | - |
| 80.80.80.0/24 | non-nh | Established | impl- | | | |
| 30.30.30.0/24 | None | No | - | | | |
| 80.80.80.0/24 | 10.10.10.2 | Established | 24328 | 0 | None | Ye |
| 80.80.80.0/24 | 20.20.20.2 | Established | impl- | | | |
| 30.30.30.0/24 | None | Yes | B | | | |
| UPSTREAM LSP : | | | | | | |
| FEC | State | Label | Req.ID | Attr | | |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes 1 |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes 1 |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes 1 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | skipped | /no 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | skipped | /no 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes 3 |
| 2.2.2.2/32 | Established | 24320 | 0 | None | Yes | /yes 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | skipped | /no 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | skipped | /no 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes 3 |
| 3.3.3.3/32 | Established | 24321 | 0 | None | Yes | /yes 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | skipped | /no 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes 3 |

| | | | | | | | |
|---------------|-------------|-----------|---|------|---------|------|---|
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes | 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | skipped | /no | 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes | 3 |
| 4.4.4.4/32 | Established | 24322 | 0 | None | Yes | /yes | 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | skipped | /no | 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes | 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes | 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | skipped | /no | 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes | 3 |
| 5.5.5.5/32 | Established | 24323 | 0 | None | Yes | /yes | 3 |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes | 1 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | skipped | /no | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | skipped | /no | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes | 3 |
| 40.40.40.0/24 | Established | 24324 | 0 | None | Yes | /yes | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | skipped | /no | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | skipped | /no | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes | 3 |
| 50.50.50.0/24 | Established | 24325 | 0 | None | Yes | /yes | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | skipped | /no | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | skipped | /no | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes | 3 |
| 60.60.60.0/24 | Established | 24326 | 0 | None | Yes | /yes | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | skipped | /no | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 3 |
| 70.70.70.0/24 | Established | 24327 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | skipped | /no | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |
| 80.80.80.0/24 | Established | 24328 | 0 | None | Yes | /yes | 3 |

PE1#show ldp fec prefix 5.5.5.5/32

LSR codes : E/N - LSR is egress/non-egress for this FEC,
 L - LSR received a label for this FEC,
 P - Primary route, B - LFA Backup route,
 R - Remote LFA Backup route,
 > - LSR will use this route for the FEC

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|------------|------|---------|-----------|-----|--------------|
| 5.5.5.5/32 | NL | 5.5.5.5 | 24329 | No | no nexthop |
| | NL | 4.4.4.4 | 24323 | No | no nexthop |
| | NLP> | 2.2.2.2 | 24325 | No | 10.10.10.2 |
| | NLB> | 3.3.3.3 | 24325 | No | 20.20.20.2 |

PE1#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|-----|---------|----------|-----------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |

```

L> 2.2.2.2/32      1      129      -      -      -      -      -
    -      N/A      06:05:28
                        1
        Yes      3      ce46/1      No      10.10.10.2      -      -
                        128
        No      24323      ce46/4      No      20.20.20.2      -      -
L> 3.3.3.3/32      2      131      -      -      -      -      -
    -      N/A      06:05:28
                        16
        Yes      24324      ce46/1      No      10.10.10.2      -      -
                        4
        No      3      ce46/4      No      20.20.20.2      -      -
L> 4.4.4.4/32      8      135      -      -      -      -      -
    -      N/A      06:05:26
                        28
        Yes      24321      ce46/1      No      10.10.10.2      -      -
                        134
        No      24324      ce46/4      No      20.20.20.2      -      -
L> 5.5.5.5/32      3      138      -      -      -      -      -
    -      N/A      06:05:28
                        9
        Yes      24325      ce46/1      No      10.10.10.2      -      -
                        7
        No      24325      ce46/4      No      20.20.20.2      -      -
L> 40.40.40.0/24   4      141      -      -      -      -      -
    -      N/A      06:05:28
                        1
        Yes      3      ce46/1      No      10.10.10.2      -      -
                        140
        No      24326      ce46/3      No      30.30.30.2      -      -
L> 50.50.50.0/24   5      143      -      -      -      -      -
    -      N/A      06:05:28
                        19
        Yes      24326      ce46/1      No      10.10.10.2      -      -
                        4
        No      3      ce46/4      No      20.20.20.2      -      -
L> 60.60.60.0/24   9      147      -      -      -      -      -
    -      N/A      06:05:26
                        22
        Yes      24327      ce46/1      No      10.10.10.2      -      -
                        146
        No      24327      ce46/4      No      20.20.20.2      -      -
L> 70.70.70.0/24   6      150      -      -      -      -      -
    -      N/A      06:05:28
                        1
        Yes      3      ce46/1      No      10.10.10.2      -      -
                        4
        No      3      ce46/4      No      20.20.20.2      -      -
L> 80.80.80.0/24   7      152      -      -      -      -      -
    -      N/A      06:05:28
                        25
        Yes      24328      ce46/1      No      10.10.10.2      -      -
                        4
        No      3      ce46/4      No      20.20.20.2      -      -
PE1#show mpls forwarding-table 5.5.5.5/32
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code  FEC      FTN-ID  Nhlfe-ID  Tunnel-ID  Pri  Out-Label  Out-
Intf  ELC      Nexthop  Algo-Num  UpTime
L> 5.5.5.5/32      3      138      -      -      -      -      -
    -      N/A      06:06:15
                        9
        Yes      24325      ce46/1      No      10.10.10.2      -      -
                        7
        No      24325      ce46/4      No      20.20.20.2      -      -

```

```

PE1#show mpls ftn-table
Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 06:06:24, UpTime: 06:06:24, LastUpdate: 05:51:47
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 refcount: 1
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 7, out intf: ce46/1, out label: 3
    Nexthop addr: 10.10.10.2          cross connect ix: 3, op code: Push

    Backup Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 128 bypass ftn-ix: 0
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 128, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/4, out label:
24323
    Nexthop addr: 20.20.20.2          cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 06:06:24, UpTime: 06:06:24, LastUpdate: 05:51:47
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 16 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 16, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24324
    Nexthop addr: 10.10.10.2          cross connect ix: 4, op code: Push

    Backup Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3
    Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 06:06:22, UpTime: 06:06:22, LastUpdate: 05:51:47
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 28 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 28, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24321
    Nexthop addr: 10.10.10.2          cross connect ix: 9, op code: Push

    Backup Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 134 bypass ftn-ix: 0
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 134, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/4, out label:
24324
    Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 06:06:24, UpTime: 06:06:24, LastUpdate: 05:51:47
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 9 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 9, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label: 24325
    Nexthop addr: 10.10.10.2          cross connect ix: 4, op code: Push

    Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 bypass ftn-ix: 0
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

```

```
Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: ce46/4, out label: 24325
Nexthop addr: 20.20.20.2          cross connect ix: 4, op code: Push
```

```
Primary FTN entry with FEC: 40.40.40.0/24, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
CreateTime: 06:06:24, UpTime: 06:06:24, LastUpdate: 05:51:47
```

```
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 refcount: 1
```

```
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 7, out intf: ce46/1, out label: 3
Nexthop addr: 10.10.10.2          cross connect ix: 3, op code: Push
```

```
Backup Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 140 bypass ftn-ix: 0
```

```
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 140, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/3, out label:
24326
```

```
Nexthop addr: 30.30.30.2          cross connect ix: 6, op code: Push
```

```
Primary FTN entry with FEC: 50.50.50.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
CreateTime: 06:06:24, UpTime: 06:06:24, LastUpdate: 05:51:47
```

```
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 19 refcount: 1
```

```
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 19, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24326
```

```
Nexthop addr: 10.10.10.2          cross connect ix: 6, op code: Push
```

```
Backup Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0
```

```
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push
```

```
Primary FTN entry with FEC: 60.60.60.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
CreateTime: 06:06:22, UpTime: 06:06:22, LastUpdate: 05:51:47
```

```
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 22 refcount: 1
```

```
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 22, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24327
```

```
Nexthop addr: 10.10.10.2          cross connect ix: 7, op code: Push
```

```
Backup Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 146 bypass ftn-ix: 0
```

```
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 146, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/4, out label:
24327
```

```
Nexthop addr: 20.20.20.2          cross connect ix: 9, op code: Push
```

```
Primary FTN entry with FEC: 70.70.70.0/24, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
CreateTime: 06:06:24, UpTime: 06:06:24, LastUpdate: 05:51:47
```

```
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 1 refcount: 1
```

```
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 7, out intf: ce46/1, out label: 3
Nexthop addr: 10.10.10.2          cross connect ix: 3, op code: Push
```

```

Backup Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2 cross connect ix: 3, op code: Push

```

Primary FTN entry with FEC: 80.80.80.0/24, id: 7, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 06:06:24, UpTime: 06:06:24, LastUpdate: 05:51:47

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 25 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 25, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:

24328

Nexthop addr: 10.10.10.2 cross connect ix: 8, op code: Push

Backup Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 4 bypass ftn-ix: 0

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 12, out intf: ce46/4, out label: 3

Nexthop addr: 20.20.20.2 cross connect ix: 3, op code: Push

PE1#show mpls ftn-table 5.5.5.5/32

Primary FTN entry with FEC: 5.5.5.5/32, id: 3, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 06:06:46, UpTime: 06:06:46, LastUpdate: 05:52:09

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 9 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 9, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label: 24325

Nexthop addr: 10.10.10.2 cross connect ix: 4, op code: Push

Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: ce46/4, out label: 24325

Nexthop addr: 20.20.20.2 cross connect ix: 4, op code: Push

PE1#show mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup

K - CLI ILM, T - MPLS-TP, s - Stitched ILM

S - SNMP, L - LDP, R - RSVP, C - CRLDP

B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT

O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI

P - SR Policy, U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|----------|----------|-----------|---------|-------------------|
| Intf/VRF | Nexthop | | pri | Algo-Num | UpTime | UPStr peers |
| L> | 4.4.4.4/32 | 2 | 24322 | 24321 | N/A | ce46/1 10.10.10.2 |
| | Yes N/A | 06:06:55 | 2 | | | |
| L> | 2.2.2.2/32 | 1 | 24320 | 3 | N/A | ce46/1 10.10.10.2 |
| | Yes N/A | 06:06:55 | 2 | | | |
| L> | 3.3.3.3/32 | 5 | 24321 | 24324 | N/A | ce46/1 10.10.10.2 |
| | Yes N/A | 06:06:55 | 2 | | | |
| L> | 50.50.50.0/24 | 6 | 24325 | 24326 | N/A | ce46/1 10.10.10.2 |
| | Yes N/A | 06:06:55 | 2 | | | |
| L> | 40.40.40.0/24 | 3 | 24324 | 3 | N/A | ce46/1 10.10.10.2 |
| | Yes N/A | 06:06:55 | 2 | | | |
| | | | 24324 | 24326 | N/A | ce46/3 30.30.30.2 |
| | No - | - | | | | |
| L> | 5.5.5.5/32 | 9 | 24323 | 24325 | N/A | ce46/1 10.10.10.2 |
| | Yes N/A | 06:06:52 | 2 | | | |
| | | | 24323 | 24325 | N/A | ce46/4 20.20.20.2 |
| | No - | - | | | | |


```

L> 70.70.70.0/24      8      24327      3      N/A      ce46/1      10.10.10.2
    Yes N/A          06:06:54  2
                                24327      3      N/A      ce46/4      20.20.20.2
    No  -            -
L> 60.60.60.0/24      4      24326      24327      N/A      ce46/1      10.10.10.2
    Yes N/A          06:06:55  2
L> 80.80.80.0/24      7      24328      24328      N/A      ce46/1      10.10.10.2
    Yes N/A          06:06:51  2
PE1#show ip ospf neighbor

Total number of full neighbors: 3
OSPF process 1 VRF(default):
Neighbor ID      Pri    State           Dead Time   Address      Interface     Instance ID
2.2.2.2          1      Full/DR         00:00:28    10.10.10.2   ce46/1        0
3.3.3.3          1      Full/DR         00:00:32    20.20.20.2   ce46/4        0
4.4.4.4          1      Full/DR         00:00:35    30.30.30.2   ce46/3        0
PE1#show ip ospf interface brief
Interface        PID    Area           Intf ID     Cost  State          Neighbors   Status
lo               1      0.0.0.0        1           1     Loopback      0           Up
ce46/1           1      0.0.0.0        10089       10    Backup        1           Up
ce46/3           1      0.0.0.0        10091       20    Backup        1           Up
ce46/4           1      0.0.0.0        10092       15    Backup        1           Up
PE1#show ip ospf database

        OSPF Router with ID (1.1.1.1) (Process ID 1 VRF default)

        Router Link States (Area 0.0.0.0)

Link ID        ADV Router    Age      Seq#          CkSum  Link count
1.1.1.1        1.1.1.1      1449     0x80000017   0xb092  4
2.2.2.2        2.2.2.2      521      0x80000012   0x9269  4
3.3.3.3        3.3.3.3      1369     0x8000001e   0x7203  5
4.4.4.4        4.4.4.4      476      0x80000018   0x3180  4
5.5.5.5        5.5.5.5      545      0x80000017   0x1705  4

        Net Link States (Area 0.0.0.0)

Link ID        ADV Router    Age      Seq#          CkSum
10.10.10.2     2.2.2.2      931      0x8000000d   0x44b5
20.20.20.2     3.3.3.3      1129     0x8000000c   0xe0f3
30.30.30.2     4.4.4.4      726      0x8000000d   0x7934
40.40.40.2     5.5.5.5      1005     0x8000000d   0x463d
50.50.50.2     5.5.5.5      865      0x8000000d   0x0f52
60.60.60.2     5.5.5.5      55       0x8000000d   0xd767
70.70.70.2     3.3.3.3      99       0x8000000d   0x0237
80.80.80.1     4.4.4.4      156      0x8000000d   0xd837

        Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID        ADV Router    Age      Seq#          CkSum  Opaque ID
1.0.0.1        1.1.1.1      939      0x8000000d   0x1803  1
1.0.0.1        2.2.2.2      651      0x8000000e   0x1af7  1
1.0.0.1        3.3.3.3      579      0x8000000e   0x1eeb  1
1.0.0.1        4.4.4.4      586      0x8000000e   0x22df  1
1.0.0.1        5.5.5.5      695      0x8000000d   0x28d2  1
1.0.0.8        2.2.2.2      1661     0x8000000c   0x6f44  8
1.0.0.10       3.3.3.3      89       0x8000000d   0x5952  10
1.0.0.14       3.3.3.3      1599     0x8000000c   0xf6de  14
1.0.0.16       3.3.3.3      519      0x8000000e   0xb16c  16
1.0.0.16       5.5.5.5      465      0x8000000e   0x9381  16
1.0.0.20       5.5.5.5      475      0x8000000e   0x11c3  20
1.0.0.24       2.2.2.2      561      0x8000000d   0xdb7b  24
1.0.0.24       4.4.4.4      476      0x8000000f   0xcfb   24
1.0.0.24       5.5.5.5      1015     0x8000000d   0x9faa  24

```

```

1.0.0.26      3.3.3.3      489      0x8000000f 0x34d4 26
1.0.0.32      2.2.2.2      511      0x8000000e 0xb64b 32
1.0.0.56      4.4.4.4      396      0x8000000e 0xa70e 56
1.0.0.62      4.4.4.4      756      0x8000000e 0x98ca 62
1.0.0.128     1.1.1.1      329      0x8000000f 0x4854 128
1.0.0.132     1.1.1.1      609      0x8000000e 0xea2c 132
1.0.0.134     1.1.1.1      869      0x8000000c 0x76e1 134
PE1#show ip route summary

```

```

-----
IP routing table name is Default-IP-Routing-Table(0)
-----

```

```

IP routing table maximum-paths      : 8
Total number of IPv4 routes         : 14
Total number of IPv4 paths          : 14
Pending routes (due to route max reached): 0
Route Source      Networks
connected         5
ospf               9
Total             14
FIB               14

```

```

ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
Total number of IPv4 ECMP paths  : 0

```

```

LFA Non ECMP statistics
-----

```

```

Total number of Routes      : 9
Total number of Primary Paths : 9
Total number of Backup Paths : 9

```

```

PE1#show ldp session

```

```

Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 5.5.5.5 | ce46/1 | Passive | OPERATIONAL | 30 | 06:07:40 |
| | 2.2.2.2 | ce46/1 | Passive | OPERATIONAL | 30 | 06:07:55 |
| | 4.4.4.4 | ce46/3 | Passive | OPERATIONAL | 30 | 06:07:55 |
| | 3.3.3.3 | ce46/4 | Passive | OPERATIONAL | 30 | 06:07:55 |

```

PE1#show ip route

```

```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

```

```

IP Route Table for VRF "default"

```

```

C      1.1.1.1/32 is directly connected, lo, installed 06:08:59, last update 06:08:59 ago
O      2.2.2.2/32 [110/11] via 10.10.10.2, ce46/1, installed 06:08:05, last update 06:08:05 ago
O      3.3.3.3/32 [110/12] via 10.10.10.2, ce46/1, installed 06:08:05, last update 06:08:05 ago
O      4.4.4.4/32 [110/13] via 10.10.10.2, ce46/1, installed 06:07:45, last update 06:07:45 ago
O      5.5.5.5/32 [110/12] via 10.10.10.2, ce46/1, installed 06:08:05, last update 06:08:05 ago
C      10.10.10.0/24 is directly connected, ce46/1, installed 06:08:56, last update 06:08:56
ago
C      20.20.20.0/24 is directly connected, ce46/4, installed 05:54:24, last update 05:54:24
ago
C      30.30.30.0/24 is directly connected, ce46/3, installed 06:08:56, last update 06:08:56
ago
O      40.40.40.0/24 [110/11] via 10.10.10.2, ce46/1, installed 06:08:05, last update 06:08:05
ago
O      50.50.50.0/24 [110/12] via 10.10.10.2, ce46/1, installed 06:07:39, last update 06:07:39
ago
O      60.60.60.0/24 [110/12] via 10.10.10.2, ce46/1, installed 06:07:42, last update 06:07:42

```

```

ago
O      70.70.70.0/24 [110/11] via 10.10.10.2, ce46/1, installed 06:08:07, last update 06:08:05
ago
O      80.80.80.0/24 [110/12] via 10.10.10.2, ce46/1, installed 06:07:45, last update 06:07:45
ago
C      127.0.0.0/8 is directly connected, lo, installed 06:09:20, last update 06:09:20 ago

Gateway of last resort is not set

```

To prohibit an interface from being used as a repair path, disable fast reroute calculation on the interface:

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface ce46/1 | Enter interface mode. |
| (config-if)#ip ospf fast-reroute per-prefix candidate disable | Disable fast reroute calculation on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit all the configurations |

Verify that the ce46/1 interface is not used for backup path calculation.

```

PE1#show ip route fast-reroute
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area ,p - stale info, E -
EVPN
       * - candidate default

IP Route Table for VRF "default"
O      3.3.3.3/32 [110/16] via 20.20.20.2, ce46/4, installed 00:23:38, last update 00:23:38 ago
       [FRR-NH] via 30.30.30.2, ce46/3

O      4.4.4.4/32 [110/17] via 20.20.20.2, ce46/4, installed 00:23:38, last update 00:23:38 ago
       [FRR-NH] via 30.30.30.2, ce46/3

O      5.5.5.5/32 [110/17] via 20.20.20.2, ce46/4, installed 00:23:38, last update 00:23:38 ago
       [FRR-NH] via 30.30.30.2, ce46/3

O      40.40.40.0/24 [110/17] via 20.20.20.2, ce46/4, installed 00:23:38, last update 00:23:38 ago
       [FRR-NH] via 30.30.30.2, ce46/3

O      60.60.60.0/24 [110/17] via 20.20.20.2, ce46/4, installed 00:23:38, last update 00:23:38 ago
       [FRR-NH] via 30.30.30.2, ce46/3

O      70.70.70.0/24 [110/16] via 20.20.20.2, ce46/4, installed 00:23:38, last update 00:23:38 ago
       [FRR-NH] via 30.30.30.2, ce46/3

O      80.80.80.0/24 [110/16] via 20.20.20.2, ce46/4, installed 00:23:38, last update 00:23:38 ago
       [FRR-NH] via 30.30.30.2, ce46/3

PE1#show ip ospf route fast-reroute
OSPF process 1:
Codes: C - connected, D - Discard, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
OSPF LFA attributes:
P - Primary, SP - Secondary-Path, LP - Link Protecting,
NP - Node Protecting, BID - Broadcast Link Protecting
DP - Downstream Protecting

O 2.2.2.2/32 [17] via 20.20.20.2, ce46/4, Area 0.0.0.0

```

```

Remote FRR path      :          via 5.5.5.5, via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [19] ,LP ,NP ,DP
O 3.3.3.3/32 [16] via 20.20.20.2, ce46/4, Area 0.0.0.0
Backup path          : via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [22] ,SP ,BID ,DP
O 4.4.4.4/32 [17] via 20.20.20.2, ce46/4, Area 0.0.0.0
Backup path          : via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [21] ,SP ,NP ,BID ,DP
O 5.5.5.5/32 [17] via 20.20.20.2, ce46/4, Area 0.0.0.0
Backup path          : via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [22] ,SP ,NP ,BID ,DP
O 30.30.30.0/24 [17] via 20.20.20.2, ce46/4, Area 0.0.0.0
Backup path          : via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [21] ,SP ,NP ,BID ,DP
O 40.40.40.0/24 [17] via 20.20.20.2, ce46/4, Area 0.0.0.0
Backup path          : via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [22] ,SP ,NP ,BID ,DP
O 50.50.50.0/24 [16] via 20.20.20.2, ce46/4, Area 0.0.0.0
Remote FRR path      :          via 5.5.5.5, via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [18] ,LP ,NP ,DP
O 60.60.60.0/24 [17] via 20.20.20.2, ce46/4, Area 0.0.0.0
Backup path          : via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [21] ,SP ,NP ,BID ,DP
O 70.70.70.0/24 [16] via 20.20.20.2, ce46/4, Area 0.0.0.0
Backup path          : via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [22] ,SP ,BID ,DP
O 80.80.80.0/24 [16] via 20.20.20.2, ce46/4, Area 0.0.0.0
Backup path          : via 30.30.30.2, ce46/3, Area 0.0.0.0
Attributes           :
Metric               : [21] ,SP ,NP ,BID ,DP

```

PE1#show ldp routes

| Prefix | Addr | NextHop | Addr | Intf | Backup | Addr | Backup | |
|---------------|-----------|------------|------------|--------|------------------|--------|--------|--------|
| Intf | Owner | CreateTime | UpdateTime | | | | | |
| 1.1.1.1/32 | | 0.0.0.0 | | lo | - | - | | |
| | connected | 07:40:43 | - | | | | | |
| 2.2.2.2/32 | | 20.20.20.2 | | ce46/4 | 5.5.5.5 | ce46/3 | ospf | 01:23: |
| 00 | 00:24:13 | | | | | | | |
| | | | | | (via 30.30.30.2) | | | |
| 3.3.3.3/32 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf | 01:23: |
| 00 | 00:24:02 | | | | | | | |
| 4.4.4.4/32 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf | 01:23: |
| 00 | 00:24:02 | | | | | | | |
| 5.5.5.5/32 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf | 01:23: |
| 00 | 00:24:02 | | | | | | | |
| 20.20.20.0/24 | | 0.0.0.0 | | ce46/4 | - | - | | |
| | connected | 07:26:08 | - | | | | | |
| 30.30.30.0/24 | | 0.0.0.0 | | ce46/3 | - | - | | |
| | connected | 07:40:40 | - | | | | | |
| 40.40.40.0/24 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf | 01:23: |
| 00 | 00:24:02 | | | | | | | |
| 50.50.50.0/24 | | 20.20.20.2 | | ce46/4 | 5.5.5.5 | ce46/3 | ospf | 01:23: |
| 00 | 00:24:13 | | | | | | | |
| | | | | | (via 30.30.30.2) | | | |
| 60.60.60.0/24 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf | 01:23: |
| 00 | 00:24:02 | | | | | | | |
| 70.70.70.0/24 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf | 01:23: |
| 00 | 00:24:02 | | | | | | | |
| 80.80.80.0/24 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf | 01:23: |
| 00 | 00:24:02 | | | | | | | |

```

PE1#show ldp fec
LSR codes      : E/N - LSR is egress/non-egress for this FEC,
                  L - LSR received a label for this FEC,
                  P - Primary route, B - LFA Backup route,
                  R - Remote LFA Backup route,
                  > - LSR will use this route for the FEC

FEC            Code  Session      Out Label  ELC  Nexthop Addr
1.1.1.1/32     NL    5.5.5.5      24333     No   connected
               NL    3.3.3.3      24328     No   connected
               NL    4.4.4.4      24329     No   connected
               NL    2.2.2.2      24329     No   connected
               E >   non-existent  none      No   connected
2.2.2.2/32     NLR> 5.5.5.5      24321     No   5.5.5.5
                                   (via 30.30.30.2,label 24323)
               NL    4.4.4.4      24321     No   no nexthop
               NLP>  3.3.3.3      24323     No   20.20.20.2
               NL    2.2.2.2      impl-null  No   no nexthop
3.3.3.3/32     NL    5.5.5.5      24322     No   no nexthop
               NLB>  4.4.4.4      24322     No   30.30.30.2
               NL    2.2.2.2      24324     No   no nexthop
               NLP>  3.3.3.3      impl-null  No   20.20.20.2
4.4.4.4/32     NL    5.5.5.5      24323     No   no nexthop
               NLB>  4.4.4.4      impl-null  No   30.30.30.2
               NL    2.2.2.2      24321     No   no nexthop
               NLP>  3.3.3.3      24324     No   20.20.20.2
5.5.5.5/32     NL    5.5.5.5      24320     No   no nexthop
               NLB>  4.4.4.4      24323     No   30.30.30.2
               NL    2.2.2.2      24325     No   no nexthop
               NLP>  3.3.3.3      24325     No   20.20.20.2
10.10.10.0/24 NL    2.2.2.2      impl-null  No   no nexthop
20.20.20.0/24  NL    5.5.5.5      24335     No   connected
               NL    4.4.4.4      24330     No   connected
               NL    2.2.2.2      24330     No   connected
               E >   non-existent  none      No   connected
               NL    3.3.3.3      impl-null  No   connected
30.30.30.0/24  NL    5.5.5.5      24326     No   connected
               NL    4.4.4.4      impl-null  No   connected
               NL    2.2.2.2      24323     No   connected
               NL    3.3.3.3      24322     No   connected
               E >   non-existent  none      No   connected
40.40.40.0/24 NL    5.5.5.5      24324     No   no nexthop
               NLB>  4.4.4.4      24326     No   30.30.30.2
               NLP>  3.3.3.3      24326     No   20.20.20.2
               NL    2.2.2.2      impl-null  No   no nexthop
50.50.50.0/24 NLR> 5.5.5.5      24325     No   5.5.5.5
                                   (via 30.30.30.2,label 24323)
               NL    4.4.4.4      24327     No   no nexthop
               NL    2.2.2.2      24326     No   no nexthop
               NLP>  3.3.3.3      impl-null  No   20.20.20.2
60.60.60.0/24 NL    5.5.5.5      24329     No   no nexthop
               NLB>  4.4.4.4      impl-null  No   30.30.30.2
               NL    2.2.2.2      24327     No   no nexthop
               NLP>  3.3.3.3      24327     No   20.20.20.2
70.70.70.0/24 NL    5.5.5.5      24327     No   no nexthop
               NLB>  4.4.4.4      24328     No   30.30.30.2
               NL    2.2.2.2      impl-null  No   no nexthop
               NLP>  3.3.3.3      impl-null  No   20.20.20.2
80.80.80.0/24 NL    5.5.5.5      24328     No   no nexthop
               NLB>  4.4.4.4      impl-null  No   30.30.30.2
               NL    2.2.2.2      24328     No   no nexthop
               NLP>  3.3.3.3      impl-null  No   20.20.20.2

PE1#show ldp downstream
FEC            Nexthop
Addr           State   Label  Req.ID  Attr          Installed  Code
Codes: P - Primary route, B - Backup route
Session peer 5.5.5.5:
80.80.80.0/24  non-

```

```

nh      Established      24328      0      None      No      -
70.70.70.0/24      non-
nh      Established      24327      0      None      No      -
60.60.60.0/24      non-
nh      Established      24329      0      None      No      -
50.50.50.0/24      30.30.30.2      Established      24325      0      None      Ye
s      B
40.40.40.0/24      non-
nh      Established      24324      0      None      No      -
30.30.30.0/24      connected      Established      24326      0      None      No
-
20.20.20.0/24      connected      Established      24335      0      None      No
-
5.5.5.5/32      non-
nh      Established      24320      0      None      No      -
4.4.4.4/32      non-
nh      Established      24323      0      None      No      -
3.3.3.3/32      non-
nh      Established      24322      0      None      No      -
2.2.2.2/32      30.30.30.2      Established      24321      0      None      Ye
s      B
1.1.1.1/32      connected      Established      24333      0      None      No
-
[Summary] total downstreams: 12
[state] depend: 0, idle: 0, resp: 0, estab: 12, inactive: 0
[label] user-label: 12, impl-null: 0, expl-null: 0, invalid: 0

Codes: P - Primary route, B - Backup route
Session peer 4.4.4.4:
20.20.20.0/24      connected      Established      24330      0      None      No
-
1.1.1.1/32      connected      Established      24329      0      None      No
-
70.70.70.0/24      30.30.30.2      Established      24328      0      None      Ye
s      B
50.50.50.0/24      non-
nh      Established      24327      0      None      No      -
40.40.40.0/24      30.30.30.2      Established      24326      0      None      Ye
s      B
30.30.30.0/24      connected      Established      impl-
null 0      None      No      -
3.3.3.3/32      30.30.30.2      Established      24322      0      None      Ye
s      B
2.2.2.2/32      non-
nh      Established      24321      0      None      No      -
80.80.80.0/24      30.30.30.2      Established      impl-
null 0      None      Yes      B
60.60.60.0/24      30.30.30.2      Established      impl-
null 0      None      Yes      B
5.5.5.5/32      30.30.30.2      Established      24323      0      None      Ye
s      B
4.4.4.4/32      30.30.30.2      Established      impl-
null 0      None      Yes      B
[Summary] total downstreams: 12
[state] depend: 0, idle: 0, resp: 0, estab: 12, inactive: 0
[label] user-label: 8, impl-null: 4, expl-null: 0, invalid: 0

Codes: P - Primary route, B - Backup route
Session peer 3.3.3.3:
1.1.1.1/32      connected      Established      24328      0      None      No
-
60.60.60.0/24      20.20.20.2      Established      24327      0      None      Ye
s      P
40.40.40.0/24      20.20.20.2      Established      24326      0      None      Ye
s      P
30.30.30.0/24      connected      Established      24322      0      None      No
-
20.20.20.0/24      connected      Established      impl-
null 0      None      No      -
4.4.4.4/32      20.20.20.2      Established      24324      0      None      Ye
s      P

```

```

2.2.2.2/32      20.20.20.2      Established      24323      0      None      Ye
s      P
80.80.80.0/24   20.20.20.2      Established      impl-
null 0      None      Yes      P
70.70.70.0/24   20.20.20.2      Established      impl-
null 0      None      Yes      P
50.50.50.0/24   20.20.20.2      Established      impl-
null 0      None      Yes      P
5.5.5.5/32      20.20.20.2      Established      24325      0      None      Ye
s      P
3.3.3.3/32      20.20.20.2      Established      impl-
null 0      None      Yes      P
[Summary] total downstreams: 12
[state] depend: 0, idle: 0, resp: 0, estab: 12, inactive: 0
[label] user-label: 7, impl-null: 5, expl-null: 0, invalid: 0

Codes: P - Primary route, B - Backup route
Session peer 2.2.2.2:
20.20.20.0/24   connected      Established      24330      0      None      No
-
1.1.1.1/32      connected      Established      24329      0      None      No
-
80.80.80.0/24   non-
nh      Established      24328      0      None      No      -
60.60.60.0/24   non-
nh      Established      24327      0      None      No      -
50.50.50.0/24   non-
nh      Established      24326      0      None      No      -
30.30.30.0/24   connected      Established      24323      0      None      No
-
10.10.10.0/24   no-route
null 0      None      No      Established      -      impl-
4.4.4.4/32      non-
nh      Established      24321      0      None      No      -
3.3.3.3/32      non-
nh      Established      24324      0      None      No      -
70.70.70.0/24   non-nh
null 0      None      No      Established      -      impl-
40.40.40.0/24   non-nh
null 0      None      No      Established      -      impl-
5.5.5.5/32      non-
nh      Established      24325      0      None      No      -
2.2.2.2/32      non-nh
null 0      None      No      Established      -      impl-
[Summary] total downstreams: 13
[state] depend: 0, idle: 0, resp: 0, estab: 13, inactive: 0
[label] user-label: 9, impl-null: 4, expl-null: 0, invalid: 0

[Summary] total sessions: 4, total downstreams: 49

PE1#show ldp lsp
DOWNSTREAM LSP :
FEC      Nextthop
Addr      State      Label      Req.ID      Attr      Code
1.1.1.1/32  connected  -      Established  24333      0      None      No
-
1.1.1.1/32  connected  -      Established  24328      0      None      No
-
1.1.1.1/32  connected  -      Established  24329      0      None      No
-
1.1.1.1/32  connected  -      Established  24329      0      None      No
-
1.1.1.1/32  connected  -      Established  none      0      None      No
-
2.2.2.2/32  30.30.30.2  Established  24321      0      None      Ye
s      B
2.2.2.2/32  non-
nh      Established  24321      0      None      No      -
2.2.2.2/32  20.20.20.2  Established  24323      0      None      Ye
s      P

```

| | | | | | | | | | |
|---------------|-------------------|-------------|-------|----|----|------|---|--|----|
| 2.2.2.2/32 | non-nh | Established | impl- | | | | | | |
| null 0 | None | No | - | | | | | | |
| 3.3.3.3/32 | non- | | | | | | | | |
| nh | Established 24322 | 0 | None | No | | | | | |
| 3.3.3.3/32 | 30.30.30.2 | Established | 24322 | 0 | No | None | - | | Ye |
| s B | | | | | | | | | |
| 3.3.3.3/32 | non- | | | | | | | | |
| nh | Established 24324 | 0 | None | No | | | | | |
| 3.3.3.3/32 | 20.20.20.2 | Established | impl- | | | | | | |
| null 0 | None | Yes | P | | | | | | |
| 4.4.4.4/32 | non- | | | | | | | | |
| nh | Established 24323 | 0 | None | No | | | | | |
| 4.4.4.4/32 | 30.30.30.2 | Established | impl- | | | | | | |
| null 0 | None | Yes | B | | | | | | |
| 4.4.4.4/32 | non- | | | | | | | | |
| nh | Established 24321 | 0 | None | No | | | | | |
| 4.4.4.4/32 | 20.20.20.2 | Established | 24324 | 0 | No | None | - | | Ye |
| s P | | | | | | | | | |
| 5.5.5.5/32 | non- | | | | | | | | |
| nh | Established 24320 | 0 | None | No | | | | | |
| 5.5.5.5/32 | 30.30.30.2 | Established | 24323 | 0 | No | None | - | | Ye |
| s B | | | | | | | | | |
| 5.5.5.5/32 | non- | | | | | | | | |
| nh | Established 24325 | 0 | None | No | | | | | |
| 5.5.5.5/32 | 20.20.20.2 | Established | 24325 | 0 | No | None | - | | Ye |
| s P | | | | | | | | | |
| 10.10.10.0/24 | no-route | Established | impl- | | | | | | |
| null 0 | None | No | - | | | | | | |
| 20.20.20.0/24 | connected | Established | 24335 | 0 | | None | | | No |
| - | | | | | | | | | |
| 20.20.20.0/24 | connected | Established | 24330 | 0 | | None | | | No |
| - | | | | | | | | | |
| 20.20.20.0/24 | connected | Established | 24330 | 0 | | None | | | No |
| - | | | | | | | | | |
| 20.20.20.0/24 | connected | Established | none | 0 | | None | | | No |
| - | | | | | | | | | |
| 20.20.20.0/24 | connected | Established | impl- | | | | | | |
| null 0 | None | No | - | | | | | | |
| 30.30.30.0/24 | connected | Established | 24326 | 0 | | None | | | No |
| - | | | | | | | | | |
| 30.30.30.0/24 | connected | Established | impl- | | | | | | |
| null 0 | None | No | - | | | | | | |
| 30.30.30.0/24 | connected | Established | 24323 | 0 | | None | | | No |
| - | | | | | | | | | |
| 30.30.30.0/24 | connected | Established | 24322 | 0 | | None | | | No |
| - | | | | | | | | | |
| 30.30.30.0/24 | connected | Established | none | 0 | | None | | | No |
| - | | | | | | | | | |
| 40.40.40.0/24 | non- | | | | | | | | |
| nh | Established 24324 | 0 | None | No | | | | | |
| 40.40.40.0/24 | 30.30.30.2 | Established | 24326 | 0 | No | None | - | | Ye |
| s B | | | | | | | | | |
| 40.40.40.0/24 | 20.20.20.2 | Established | 24326 | 0 | | None | | | Ye |
| s P | | | | | | | | | |
| 40.40.40.0/24 | non-nh | Established | impl- | | | | | | |
| null 0 | None | No | - | | | | | | |
| 50.50.50.0/24 | 30.30.30.2 | Established | 24325 | 0 | | None | | | Ye |
| s B | | | | | | | | | |
| 50.50.50.0/24 | non- | | | | | | | | |
| nh | Established 24327 | 0 | None | No | | | | | |
| 50.50.50.0/24 | non- | | | | | | | | |
| nh | Established 24326 | 0 | None | No | | | | | |
| 50.50.50.0/24 | 20.20.20.2 | Established | impl- | | | | | | |
| null 0 | None | Yes | P | | | | | | |
| 60.60.60.0/24 | non- | | | | | | | | |
| nh | Established 24329 | 0 | None | No | | | | | |
| 60.60.60.0/24 | 30.30.30.2 | Established | impl- | | | | | | |
| null 0 | None | Yes | B | | | | | | |
| 60.60.60.0/24 | non- | | | | | | | | |
| nh | Established 24327 | 0 | None | No | | | | | |
| 60.60.60.0/24 | 20.20.20.2 | Established | 24327 | 0 | No | None | - | | Ye |

| | | | | | | | | | |
|----------------|-------------|-------------|--------|------|---------|------|---|--|----|
| s | P | | | | | | | | |
| 70.70.70.0/24 | non- | | | | | | | | |
| nh | Established | 24327 | 0 | None | No | - | | | |
| 70.70.70.0/24 | 30.30.30.2 | Established | 24328 | 0 | None | | | | Ye |
| s | B | | | | | | | | |
| 70.70.70.0/24 | non-nh | Established | impl- | | | | | | |
| null 0 | None | No | - | | | | | | |
| 70.70.70.0/24 | 20.20.20.2 | Established | impl- | | | | | | |
| null 0 | None | Yes | P | | | | | | |
| 80.80.80.0/24 | non- | | | | | | | | |
| nh | Established | 24328 | 0 | None | No | - | | | |
| 80.80.80.0/24 | 30.30.30.2 | Established | impl- | | | | | | |
| null 0 | None | Yes | B | | | | | | |
| 80.80.80.0/24 | non- | | | | | | | | |
| nh | Established | 24328 | 0 | None | No | - | | | |
| 80.80.80.0/24 | 20.20.20.2 | Established | impl- | | | | | | |
| null 0 | None | Yes | P | | | | | | |
| UPSTREAM LSP : | | | | | | | | | |
| FEC | State | Label | Req.ID | Attr | | | | | |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes | 1 | | |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /no | 1 | | |
| 1.1.1.1/32 | Established | impl-null | 0 | None | No | /yes | 1 | | |
| 1.1.1.1/32 | Established | 24320 | 0 | None | Yes | /yes | 1 | | |
| 2.2.2.2/32 | Established | 24333 | 0 | None | Yes | /yes | 4 | | |
| 2.2.2.2/32 | Established | 24333 | 0 | None | Yes | /yes | 4 | | |
| 2.2.2.2/32 | Established | 24333 | 0 | None | skipped | /no | 4 | | |
| 2.2.2.2/32 | Established | 24333 | 0 | None | skipped | /no | 4 | | |
| 2.2.2.2/32 | Established | 24333 | 0 | None | skipped | /no | 4 | | |
| 2.2.2.2/32 | Established | 24333 | 0 | None | Yes | /yes | 4 | | |
| 2.2.2.2/32 | Established | 24333 | 0 | None | skipped | /no | 4 | | |
| 2.2.2.2/32 | Established | 24333 | 0 | None | Yes | /yes | 4 | | |
| 3.3.3.3/32 | Established | 24334 | 0 | None | Yes | /yes | 4 | | |
| 3.3.3.3/32 | Established | 24334 | 0 | None | Yes | /yes | 4 | | |
| 3.3.3.3/32 | Established | 24334 | 0 | None | skipped | /no | 4 | | |
| 3.3.3.3/32 | Established | 24334 | 0 | None | skipped | /no | 4 | | |
| 3.3.3.3/32 | Established | 24334 | 0 | None | skipped | /no | 4 | | |
| 3.3.3.3/32 | Established | 24334 | 0 | None | Yes | /yes | 4 | | |
| 3.3.3.3/32 | Established | 24334 | 0 | None | skipped | /no | 4 | | |
| 3.3.3.3/32 | Established | 24334 | 0 | None | Yes | /yes | 4 | | |
| 4.4.4.4/32 | Established | 24335 | 0 | None | Yes | /yes | 4 | | |
| 4.4.4.4/32 | Established | 24335 | 0 | None | skipped | /no | 4 | | |
| 4.4.4.4/32 | Established | 24335 | 0 | None | skipped | /no | 4 | | |
| 4.4.4.4/32 | Established | 24335 | 0 | None | skipped | /no | 4 | | |
| 4.4.4.4/32 | Established | 24335 | 0 | None | Yes | /yes | 4 | | |
| 4.4.4.4/32 | Established | 24335 | 0 | None | skipped | /no | 4 | | |
| 4.4.4.4/32 | Established | 24335 | 0 | None | Yes | /yes | 4 | | |
| 5.5.5.5/32 | Established | 24336 | 0 | None | Yes | /yes | 4 | | |
| 5.5.5.5/32 | Established | 24336 | 0 | None | Yes | /yes | 4 | | |
| 5.5.5.5/32 | Established | 24336 | 0 | None | skipped | /no | 4 | | |
| 5.5.5.5/32 | Established | 24336 | 0 | None | skipped | /no | 4 | | |
| 5.5.5.5/32 | Established | 24336 | 0 | None | skipped | /no | 4 | | |
| 5.5.5.5/32 | Established | 24336 | 0 | None | Yes | /yes | 4 | | |
| 5.5.5.5/32 | Established | 24336 | 0 | None | skipped | /no | 4 | | |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /no | 1 | | |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes | 1 | | |
| 20.20.20.0/24 | Established | 24321 | 0 | None | Yes | /yes | 1 | | |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes | 1 | | |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /no | 1 | | |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes | 1 | | |
| 30.30.30.0/24 | Established | 24322 | 0 | None | Yes | /yes | 1 | | |
| 40.40.40.0/24 | Established | 24337 | 0 | None | Yes | /yes | 4 | | |
| 40.40.40.0/24 | Established | 24337 | 0 | None | Yes | /yes | 4 | | |
| 40.40.40.0/24 | Established | 24337 | 0 | None | skipped | /no | 4 | | |
| 40.40.40.0/24 | Established | 24337 | 0 | None | skipped | /no | 4 | | |
| 40.40.40.0/24 | Established | 24337 | 0 | None | Yes | /yes | 4 | | |
| 40.40.40.0/24 | Established | 24337 | 0 | None | skipped | /no | 4 | | |

| | | | | | | | |
|---------------|-------------|-------|---|------|---------|------|---|
| 40.40.40.0/24 | Established | 24337 | 0 | None | skipped | /no | 4 |
| 40.40.40.0/24 | Established | 24337 | 0 | None | Yes | /yes | 4 |
| 50.50.50.0/24 | Established | 24338 | 0 | None | Yes | /yes | 4 |
| 50.50.50.0/24 | Established | 24338 | 0 | None | Yes | /yes | 4 |
| 50.50.50.0/24 | Established | 24338 | 0 | None | skipped | /no | 4 |
| 50.50.50.0/24 | Established | 24338 | 0 | None | skipped | /no | 4 |
| 50.50.50.0/24 | Established | 24338 | 0 | None | skipped | /no | 4 |
| 50.50.50.0/24 | Established | 24338 | 0 | None | Yes | /yes | 4 |
| 50.50.50.0/24 | Established | 24338 | 0 | None | skipped | /no | 4 |
| 50.50.50.0/24 | Established | 24338 | 0 | None | Yes | /yes | 4 |
| 60.60.60.0/24 | Established | 24339 | 0 | None | Yes | /yes | 4 |
| 60.60.60.0/24 | Established | 24339 | 0 | None | Yes | /yes | 4 |
| 60.60.60.0/24 | Established | 24339 | 0 | None | skipped | /no | 4 |
| 60.60.60.0/24 | Established | 24339 | 0 | None | skipped | /no | 4 |
| 60.60.60.0/24 | Established | 24339 | 0 | None | skipped | /no | 4 |
| 60.60.60.0/24 | Established | 24339 | 0 | None | Yes | /yes | 4 |
| 60.60.60.0/24 | Established | 24339 | 0 | None | skipped | /no | 4 |
| 60.60.60.0/24 | Established | 24339 | 0 | None | Yes | /yes | 4 |
| 70.70.70.0/24 | Established | 24340 | 0 | None | Yes | /yes | 4 |
| 70.70.70.0/24 | Established | 24340 | 0 | None | Yes | /yes | 4 |
| 70.70.70.0/24 | Established | 24340 | 0 | None | skipped | /no | 4 |
| 70.70.70.0/24 | Established | 24340 | 0 | None | skipped | /no | 4 |
| 70.70.70.0/24 | Established | 24340 | 0 | None | skipped | /no | 4 |
| 70.70.70.0/24 | Established | 24340 | 0 | None | Yes | /yes | 4 |
| 70.70.70.0/24 | Established | 24340 | 0 | None | skipped | /no | 4 |
| 70.70.70.0/24 | Established | 24340 | 0 | None | Yes | /yes | 4 |
| 80.80.80.0/24 | Established | 24341 | 0 | None | Yes | /yes | 4 |
| 80.80.80.0/24 | Established | 24341 | 0 | None | Yes | /yes | 4 |
| 80.80.80.0/24 | Established | 24341 | 0 | None | skipped | /no | 4 |
| 80.80.80.0/24 | Established | 24341 | 0 | None | skipped | /no | 4 |
| 80.80.80.0/24 | Established | 24341 | 0 | None | skipped | /no | 4 |
| 80.80.80.0/24 | Established | 24341 | 0 | None | Yes | /yes | 4 |
| 80.80.80.0/24 | Established | 24341 | 0 | None | skipped | /no | 4 |
| 80.80.80.0/24 | Established | 24341 | 0 | None | Yes | /yes | 4 |

```
PE1#show ldp fec prefix 5.5.5.5/32
```

```
LSR codes : E/N - LSR is egress/non-egress for this FEC,
            L - LSR received a label for this FEC,
            P - Primary route, B - LFA Backup route,
            R - Remote LFA Backup route,
            > - LSR will use this route for the FEC
```

| FEC | Code | Session | Out Label | ELC | Nexthop Addr |
|------------|------|---------|-----------|-----|--------------|
| 5.5.5.5/32 | NL | 5.5.5.5 | 24320 | No | no nexthop |
| | NLB> | 4.4.4.4 | 24323 | No | 30.30.30.2 |
| | NL | 2.2.2.2 | 24325 | No | no nexthop |
| | NLP> | 3.3.3.3 | 24325 | No | 20.20.20.2 |

```
PE1#show mpls forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
        B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
        (m) - FTN mapped over multipath transport, (e) - FTN is ECMP
```

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|------------|---------|----------|-----------|------------|-----------|--------------------------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 2.2.2.2/32 | 1 | 59 | - | - | - | - |
| | - | N/A | 01:24:14 | | | | |
| | | | 21 | - | | | |
| | Yes | 24323 | ce46/4 | No | 20.20.20.2 | - | - |
| | | | 58 | - | | | |
| | No | 24321 | ce46/3 | No | 5.5.5.5 | - | - |
| | | | | | | | (via |
| | | | | | | | 30.30.30.2, label 24323) |
| L> | 3.3.3.3/32 | 2 | 63 | - | - | - | - |
| | - | N/A | 01:24:14 | | | | |
| | | | 4 | - | | | |
| | Yes | 3 | ce46/4 | No | 20.20.20.2 | - | - |

```

        62
L>      No      24322      ce46/3      No      30.30.30.2      -      -
      4.4.4.4/32      3      65      -      -      -
      -      N/A      01:24:14
        30
      Yes      24324      ce46/4      No      20.20.20.2      -      -
        10
L>      No      3      ce46/3      No      30.30.30.2      -      -
      5.5.5.5/32      4      69      -      -      -
      -      N/A      01:24:14
        7
      Yes      24325      ce46/4      No      20.20.20.2      -      -
        68
L>      No      24323      ce46/3      No      30.30.30.2      -      -
      40.40.40.0/24      5      75      -      -      -
      -      N/A      01:24:14
        53
      Yes      24326      ce46/4      No      20.20.20.2      -      -
        74
L>      No      24326      ce46/3      No      30.30.30.2      -      -
      50.50.50.0/24      6      61      -      -      -
      -      N/A      01:24:14
        4
      Yes      3      ce46/4      No      20.20.20.2      -      -
        60
      No      24325      ce46/3      No      5.5.5.5      -      -
30.30.30.2, label 24323)
L>      60.60.60.0/24      7      78      -      -      -
      -      N/A      01:24:14
        43
      Yes      24327      ce46/4      No      20.20.20.2      -      -
        10
L>      No      3      ce46/3      No      30.30.30.2      -      -
      70.70.70.0/24      8      82      -      -      -
      -      N/A      01:24:14
        4
      Yes      3      ce46/4      No      20.20.20.2      -      -
        81
L>      No      24328      ce46/3      No      30.30.30.2      -      -
      80.80.80.0/24      9      84      -      -      -
      -      N/A      01:24:14
        4
      Yes      3      ce46/4      No      20.20.20.2      -      -
        10
      No      3      ce46/3      No      30.30.30.2      -      -

PE1#show mpls forwarding-table 5.5.5.5/32
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code   FEC           FTN-ID   Nhlfe-ID  Tunnel-ID   Pri   Out-Label   Out-
Intf   ELC           NextHop   Algo-Num   UpTime
L>     5.5.5.5/32      4        69        -          -      -          -
      -      N/A      01:24:27
        7
      Yes      24325      ce46/4      No      20.20.20.2      -      -
        68
      No      24323      ce46/3      No      30.30.30.2      -      -

PE1#show mpls ftn-table
Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:24:37, UpTime: 01:24:37, LastUpdate: 00:25:40
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0

```

```

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 21 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 21, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24323
  Nexthop addr: 20.20.20.2          cross connect ix: 1, op code: Push

  Backup Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 58 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 58, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/3, out label:
24321
  Nexthop addr: 5.5.5.5            cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:24:37, UpTime: 01:24:37, LastUpdate: 00:25:40
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 15, out intf: ce46/4, out label: 3
  Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

  Backup Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 62 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 62, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/3, out label:
24322
  Nexthop addr: 30.30.30.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:24:37, UpTime: 01:24:37, LastUpdate: 00:25:40
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 30 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 30, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24324
  Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

  Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 10 bypass ftn-ix: 0
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 10, owner: N/A, Stale: NO, refcount: 6, out intf: ce46/3, out label: 3
  Nexthop addr: 30.30.30.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:24:37, UpTime: 01:24:37, LastUpdate: 00:25:40
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: ce46/4, out label: 24325
  Nexthop addr: 20.20.20.2          cross connect ix: 4, op code: Push

  Backup Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 68 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 68, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/3, out label:
24323
  Nexthop addr: 30.30.30.2          cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 40.40.40.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed

```

```
CreateTime: 01:24:37, UpTime: 01:24:37, LastUpdate: 00:25:40
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 53 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 53, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24326
Nexthop addr: 20.20.20.2          cross connect ix: 11, op code: Push

Backup Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 74 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 74, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/3, out label:
24326
Nexthop addr: 30.30.30.2          cross connect ix: 8, op code: Push

Primary FTN entry with FEC: 50.50.50.0/24, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:24:37, UpTime: 01:24:37, LastUpdate: 00:25:40
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 15, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

Backup Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 60 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 60, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/3, out label:
24325
Nexthop addr: 5.5.5.5            cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 60.60.60.0/24, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:24:37, UpTime: 01:24:37, LastUpdate: 00:25:40
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 43 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 43, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24327
Nexthop addr: 20.20.20.2          cross connect ix: 8, op code: Push

Backup Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 10 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 10, owner: N/A, Stale: NO, refcount: 6, out intf: ce46/3, out label: 3
Nexthop addr: 30.30.30.2          cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 70.70.70.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:24:37, UpTime: 01:24:37, LastUpdate: 00:25:40
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 15, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 3, op code: Push

Backup Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 81 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 81, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/3, out label:
24328
Nexthop addr: 30.30.30.2          cross connect ix: 10, op code: Push
```

Primary FTN entry with FEC: 80.80.80.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:24:37, UpTime: 01:24:37, LastUpdate: 00:25:40

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 15, out intf: ce46/4, out label: 3

Nexthop addr: 20.20.20.2 cross connect ix: 3, op code: Push

Backup Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 10 bypass ftn-ix: 0

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 10, owner: N/A, Stale: NO, refcount: 6, out intf: ce46/3, out label: 3

Nexthop addr: 30.30.30.2 cross connect ix: 3, op code: Push

PE1#show mpls ftn-table 5.5.5.5/32

Primary FTN entry with FEC: 5.5.5.5/32, id: 4, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:24:53, UpTime: 01:24:53, LastUpdate: 00:25:56

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: ce46/4, out label: 24325

Nexthop addr: 20.20.20.2 cross connect ix: 4, op code: Push

Backup Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 68 bypass ftn-ix: 0

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 68, owner: LDP, Stale: NO, refcount: 1, out intf: ce46/3, out label:

24323

Nexthop addr: 30.30.30.2 cross connect ix: 5, op code: Push

PE1#show mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup

K - CLI ILM, T - MPLS-TP, s - Stitched ILM

S - SNMP, L - LDP, R - RSVP, C - CRLDP

B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT

O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI

P - SR Policy, U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- | |
|--------------------------|---------------|----------|----------|-----------|---------|--------|------------|
| Intf/VRF | Nexthop | | pri | Algo-Num | UpTime | UPStr | peers |
| L> 3.3.3.3/32 | Yes N/A | 2 | 24334 | 3 | N/A | ce46/4 | 20.20.20.2 |
| | Yes N/A | 00:26:20 | 2 | | | | |
| L> 20.20.20.0/24 | Yes N/A | 10 | 24321 | Nolabel | N/A | N/A | 127.0.0.1 |
| | Yes N/A | 00:26:10 | 1 | | | | |
| L> 1.1.1.1/32 | Yes N/A | 5 | 24320 | Nolabel | N/A | N/A | 127.0.0.1 |
| | Yes N/A | 00:26:10 | 1 | | | | |
| L> 2.2.2.2/32 | Yes N/A | 1 | 24333 | 24323 | N/A | ce46/4 | 20.20.20.2 |
| | Yes N/A | 00:26:20 | 2 | | | | |
| | | | 24333 | 24321 | N/A | ce46/3 | 5.5.5.5 |
| | No - | - | | | | | (via |
| 30.30.30.2, label 24323) | | | | | | | |
| L> 30.30.30.0/24 | Yes N/A | 11 | 24322 | Nolabel | N/A | N/A | 127.0.0.1 |
| | Yes N/A | 00:26:10 | 1 | | | | |
| L> 50.50.50.0/24 | Yes N/A | 6 | 24338 | 3 | N/A | ce46/4 | 20.20.20.2 |
| | Yes N/A | 00:26:20 | 2 | | | | |
| | | | 24338 | 24325 | N/A | ce46/3 | 5.5.5.5 |
| | No - | - | | | | | (via |
| 30.30.30.2, label 24323) | | | | | | | |
| L> 5.5.5.5/32 | Yes N/A | 4 | 24336 | 24325 | N/A | ce46/4 | 20.20.20.2 |
| | Yes N/A | 00:26:20 | 2 | | | | |
| L> 4.4.4.4/32 | | 3 | 24335 | 24324 | N/A | ce46/4 | 20.20.20.2 |

```

      Yes  N/A      00:26:20  2
L> 40.40.40.0/24    18      24337      24326      N/A      ce46/4      20.20.20.2
      Yes  N/A      01:25:07  2
L> 70.70.70.0/24    8      24340      3          N/A      ce46/4      20.20.20.2
      Yes  N/A      00:26:20  2
L> 60.60.60.0/24    7      24339      24327      N/A      ce46/4      20.20.20.2
      Yes  N/A      00:26:20  2
L> 80.80.80.0/24    9      24341      3          N/A      ce46/4      20.20.20.2
      Yes  N/A      00:26:20  2

PE1#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID    Pri    State           Dead Time   Address        Interface      Instance ID
3.3.3.3        1      Full/DR         00:00:32    20.20.20.2     ce46/4         0
4.4.4.4        1      Full/DR         00:00:38    30.30.30.2     ce46/3         0
PE1#show ip ospf int brief
Interface      PID    Area           Intf ID     Cost    State          Neighbors    Status
lo             1      0.0.0.0        1           1      Loopback       0            Up
ce46/1         1      0.0.0.0        10089       10      Down           0            Down
ce46/3         1      0.0.0.0        10091       20      Backup         1            Up
ce46/4         1      0.0.0.0        10092       15      Backup         1            Up

PE1#show ip ospf database

        OSPF Router with ID (1.1.1.1) (Process ID 1 VRF default)

        Router Link States (Area 0.0.0.0)

Link ID        ADV Router      Age      Seq#           CkSum  Link count
1.1.1.1        1.1.1.1        1126    0x80000008    0x9d0d  3
2.2.2.2        2.2.2.2        1128    0x8000001a    0xe261  3
3.3.3.3        3.3.3.3        1545    0x80000025    0x640a  5
4.4.4.4        4.4.4.4        1582    0x8000001e    0x2586  4
5.5.5.5        5.5.5.5        881     0x8000001a    0x1108  4

        Net Link States (Area 0.0.0.0)

Link ID        ADV Router      Age      Seq#           CkSum
20.20.20.2     3.3.3.3        1125    0x80000003    0xf2ea
30.30.30.2     4.4.4.4        1602    0x80000011    0x7138
40.40.40.2     5.5.5.5        1341    0x80000010    0x4040
50.50.50.2     5.5.5.5        1231    0x80000010    0x0955
60.60.60.2     5.5.5.5        511     0x80000010    0xd16a
70.70.70.2     3.3.3.3        405     0x80000010    0xfb3a
80.80.80.1     4.4.4.4        632     0x80000010    0xd23a

        Area-Local Opaque-LSA (Area 0.0.0.0)

Link ID        ADV Router      Age      Seq#           CkSum  Opaque ID
1.0.0.1        1.1.1.1        1595    0x80000004    0x2af9  1
1.0.0.1        2.2.2.2        1058    0x80000011    0x14fa  1
1.0.0.1        3.3.3.3        1055    0x80000011    0x18ee  1
1.0.0.1        4.4.4.4        902     0x80000011    0x1ce2  1
1.0.0.1        5.5.5.5        1001    0x80000010    0x22d5  1
1.0.0.8        2.2.2.2        298     0x80000010    0x6748  8
1.0.0.10       3.3.3.3        495     0x80000010    0x5355  10
1.0.0.14       3.3.3.3        1085    0x80000003    0x09d5  14
1.0.0.16       3.3.3.3        775     0x80000011    0xab6f  16
1.0.0.16       5.5.5.5        691     0x80000011    0x8d84  16
1.0.0.20       5.5.5.5        891     0x80000011    0x0bc6  20
1.0.0.24       2.2.2.2        938     0x80000010    0xd57e  24
1.0.0.24       4.4.4.4        792     0x80000012    0x06fe  24
1.0.0.24       5.5.5.5        1311    0x80000010    0x99ad  24

```

| | | | | | |
|-----------|---------|------|------------|--------|-----|
| 1.0.0.26 | 3.3.3.3 | 795 | 0x80000012 | 0x2ed7 | 26 |
| 1.0.0.56 | 4.4.4.4 | 682 | 0x80000011 | 0xa111 | 56 |
| 1.0.0.62 | 4.4.4.4 | 172 | 0x80000004 | 0xacc0 | 62 |
| 1.0.0.132 | 1.1.1.1 | 1605 | 0x80000004 | 0xfe22 | 132 |
| 1.0.0.134 | 1.1.1.1 | 1565 | 0x80000004 | 0x86d9 | 134 |

```
PE1#show ip route summary
```

```
-----
IP routing table name is Default-IP-Routing-Table(0)
-----
```

```
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 13
Total number of IPv4 paths       : 13
Pending routes (due to route max reached): 0
Route Source   Networks
connected      4
ospf            9
Total          13
FIB             13
```

```
ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
Total number of IPv4 ECMP paths  : 0
```

```
LFA Non ECMP statistics
-----
```

```
Total number of Routes      : 7
Total number of Primary Paths : 7
Total number of Backup Paths : 7
```

```
PE1#show ip interface brief
```

```
'*' - address is assigned by dhcp client
```

| Interface | IP-Address | Admin-Status | Link-Status |
|-----------|------------|--------------|-------------|
| cd48 | unassigned | up | down |
| cd49 | unassigned | up | down |
| cd50 | unassigned | up | down |
| cd51 | unassigned | up | down |
| cd52 | unassigned | up | down |
| cd53 | unassigned | up | down |
| cd54 | unassigned | up | down |
| cd55 | unassigned | up | down |
| ce1 | unassigned | up | down |
| ce2 | unassigned | up | down |
| ce3 | unassigned | up | down |
| ce4 | unassigned | up | down |
| ce5 | unassigned | up | down |
| ce6 | unassigned | up | down |
| ce7 | unassigned | up | down |
| ce8 | unassigned | up | down |
| ce9 | unassigned | up | down |
| ce10 | unassigned | up | down |
| ce11 | unassigned | up | down |
| ce12 | unassigned | up | down |
| ce13 | unassigned | up | down |
| ce14 | unassigned | up | down |
| ce15 | unassigned | up | down |
| ce16 | unassigned | up | down |
| ce17 | unassigned | up | down |
| ce18 | unassigned | up | down |
| ce19 | unassigned | up | down |
| ce20 | unassigned | up | down |
| ce21 | unassigned | up | down |
| ce22 | unassigned | up | down |
| ce23 | unassigned | up | down |
| ce24 | unassigned | up | down |


```

ce25          unassigned    up          down
ce26          unassigned    up          down
ce27          unassigned    up          down
ce28          unassigned    up          down
ce29          unassigned    up          down
ce30          unassigned    up          down
ce31          unassigned    up          down
ce32          unassigned    up          down
ce33          unassigned    up          down
ce34          unassigned    up          down
ce35          unassigned    up          down
ce36          unassigned    up          down
ce37          unassigned    up          down
ce38          unassigned    up          down
ce39          unassigned    up          down
ce40          unassigned    up          down
ce41          unassigned    up          down
ce42          unassigned    up          down
ce43          unassigned    up          down
ce44          unassigned    up          up
ce45          unassigned    up          down
ce46/1        10.10.10.1    administratively down down
ce46/2        unassigned    up          up
ce46/3        30.30.30.1    up          up
ce46/4        20.20.20.1    up          up
ce47          unassigned    up          down
eth0          *10.16.119.123    up          up
lo            127.0.0.1      up          up
lo.management 127.0.0.1      up          up
xe0           unassigned    up          down
xe1           unassigned    up          up
xe2           unassigned    up          down
xe3           unassigned    up          down

```

PE1#show ldp session

```

Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 5.5.5.5 | ce46/3 | Passive | OPERATIONAL | 30 | 01:26:21 |
| | 4.4.4.4 | ce46/3 | Passive | OPERATIONAL | 30 | 07:43:25 |
| | 3.3.3.3 | ce46/4 | Passive | OPERATIONAL | 30 | 07:43:25 |
| | 2.2.2.2 | ce46/4 | Passive | OPERATIONAL | 30 | 07:43:25 |

PE1#show ip route

```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

```

IP Route Table for VRF "default"

```

C      1.1.1.1/32 is directly connected, lo, installed 07:44:20, last update 07:44:20 ago
O      2.2.2.2/32 [110/17] via 20.20.20.2, ce46/4, installed 00:27:50, last update 00:27:50 ago
O      3.3.3.3/32 [110/16] via 20.20.20.2, ce46/4, installed 00:27:50, last update 00:27:50 ago
O      4.4.4.4/32 [110/17] via 20.20.20.2, ce46/4, installed 00:27:50, last update 00:27:50 ago
O      5.5.5.5/32 [110/17] via 20.20.20.2, ce46/4, installed 00:27:50, last update 00:27:50 ago
C      20.20.20.0/24 is directly connected, ce46/4, installed 07:29:45, last update 07:29:45 ago
C      30.30.30.0/24 is directly connected, ce46/3, installed 07:44:17, last update 07:44:17 ago
O      40.40.40.0/24 [110/17] via 20.20.20.2, ce46/4, installed 00:27:50, last update 00:27:50 ago
O      50.50.50.0/24 [110/16] via 20.20.20.2, ce46/4, installed 00:27:50, last update 00:27:50 ago

```

```

ago
O      60.60.60.0/24 [110/17] via 20.20.20.2, ce46/4, installed 00:27:50, last update 00:27:50
ago
O      70.70.70.0/24 [110/16] via 20.20.20.2, ce46/4, installed 00:27:50, last update 00:27:50
ago
O      80.80.80.0/24 [110/16] via 20.20.20.2, ce46/4, installed 00:27:50, last update 00:27:50
ago
C      127.0.0.0/8 is directly connected, lo, installed 07:44:41, last update 07:44:41 ago

```

Gateway of last resort is not set

PE1#show ldp routes


| Prefix | Addr | Nexthop | Addr | Intf | Backup | Addr | Backup |
|---------------|-----------|------------|------------|--------|------------------|--------|-------------|
| Intf | Owner | CreateTime | UpdateTime | | | | |
| 1.1.1.1/32 | | 0.0.0.0 | | lo | - | | - |
| | connected | 07:44:32 | - | | | | |
| 2.2.2.2/32 | | 20.20.20.2 | | ce46/4 | 5.5.5.5 | ce46/3 | ospf 01:26: |
| 49 | 00:28:02 | | | | | | |
| | | | | | (via 30.30.30.2) | | |
| 3.3.3.3/32 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf 01:26: |
| 49 | 00:27:51 | | | | | | |
| 4.4.4.4/32 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf 01:26: |
| 49 | 00:27:51 | | | | | | |
| 5.5.5.5/32 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf 01:26: |
| 49 | 00:27:51 | | | | | | |
| 20.20.20.0/24 | | 0.0.0.0 | | ce46/4 | - | | - |
| | connected | 07:29:57 | - | | | | |
| 30.30.30.0/24 | | 0.0.0.0 | | ce46/3 | - | | - |
| | connected | 07:44:29 | - | | | | |
| 40.40.40.0/24 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf 01:26: |
| 49 | 00:27:51 | | | | | | |
| 50.50.50.0/24 | | 20.20.20.2 | | ce46/4 | 5.5.5.5 | ce46/3 | ospf 01:26: |
| 49 | 00:28:02 | | | | | | |
| | | | | | (via 30.30.30.2) | | |
| 60.60.60.0/24 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf 01:26: |
| 49 | 00:27:51 | | | | | | |
| 70.70.70.0/24 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf 01:26: |
| 49 | 00:27:51 | | | | | | |
| 80.80.80.0/24 | | 20.20.20.2 | | ce46/4 | 30.30.30.2 | ce46/3 | ospf 01:26: |
| 49 | 00:27:51 | | | | | | |

LDP-FRR with ISIS as IGP Configuration

Below are the configurations and validations involving NSM, ISIS, LDP before configuring fast-reroute for IGP and LDP.

PE1


| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1 |
| (config-router)#transport-address ipv4 1.1.1.1 0 | Configure the transport address for IPv4 (for IPv6 use |

| | |
|---|--|
| | <p>ipv6) to be used for a TCP session over which LDP will run.</p> <div>  Note: It is preferable to use the loopback address as the transport address. </div> |
| <code>(config-router)#targeted-peer ipv4 5.5.5.5</code> | Configure targeted peer. |
| <code>(config-router-targeted-peer)#exit</code> | Exit-targeted-peer-mode |
| <code>(config-router)#exit</code> | Exit router mode |
| <code>(config)#interface ce46/1</code> | Enter interface mode. |
| <code>(config-if)#ip address 10.10.10.1/24</code> | Configure IPv4 address for ce46/1. |
| <code>(config-if)#ip router isis 1</code> | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface ce46/1. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on ce46/1. |
| <code>(config-if)#isis wide-metric 20 level-1</code> | Assign isis wide-metric to the interface |
| <code>(config-if)#mpls ldp-igp sync isis level-1</code> | <p>Configure LDP-IGP Synchronization for interface ce46/1 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command.</p> <p>IS-IS: This command is part of ISIS Process.</p> <p>Default: Mandatory configuration. No default option.</p> |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface ce46/4</code> | Enter interface mode. |
| <code>(config-if)#ip address 20.20.20.1/24</code> | Configure IPv4 address for ce46/4. |
| <code>(config-if)#ip router isis 1</code> | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface ce46/4. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on ce46/4. |
| <code>(config-if)#isis wide-metric 15 level-1</code> | Assign isis wide-metric to the interface |
| <code>(config-if)#mpls ldp-igp sync isis level-1</code> | <p>Configure LDP-IGP Synchronization for interface ce46/4 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command.</p> |

| | |
|---|---|
| | IS-IS: This command is part of ISIS Process. Default: Mandatory configuration. No default option. |
| (config-if)#exit | Exit interface mode |
| (config)#interface ce46/3 | Enter interface mode. |
| (config-if)#ip address 30.30.30.1/24 | Configure IPv4 address for ce46/3. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface ce46/3. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on ce46/3. |
| (config-if)#isis wide-metric 60 level-1 | Assign isis wide-metric to the interface |
| (config-if)#mpls ldp-igp sync isis level-1 | Configure LDP-IGP Synchronization for interface ce46/3 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command. IS-IS: This command is part of ISIS Process. Default: Mandatory configuration. No default option. |
| (config-if)#exit | Exit interface mode |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 with instance 1 |
| (config-router)#net 49.0001.0000.0000.0001.00 | Establish a Network Entity Title for this instance, specifying the area address and the system ID |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config-router)#is-type level-1 | Configure instance as level-1-only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1. |
| (config-router)#bfd all-interfaces | Enable BFD for ISIS on all interfaces |
| (config-router)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config-if)#exit | Exit interface mode. |
| (config)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval |
| (config)#commit | Commit all the configurations |

P1


| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to |

| | |
|--|--|
| | 2.2.2.2/32 |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| (config-router)#transport-address ipv4 2.2.2.2 0 | <p>Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run.</p> <div>  <p>Note: It is preferable to use the loopback address as the transport address.</p> </div> |
| (config-router)#exit | Exit router mode |
| (config)#interface xe12 | Enter interface mode. |
| (config-if)#ip address 10.10.10.2/24 | Configure IPv4 address for xe12. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe12. |
| (config-if)#mpls ldp-igp sync isis level-1 | <p>Configure LDP-IGP Synchronization for interface xe12 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command.</p> <p>IS-IS: This command is part of ISIS Process.</p> <p>Default: Mandatory configuration. No default option.</p> |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe8 | Enter interface mode. |
| (config-if)#ip address 40.40.40.1/24 | Configure IPv4 address for xe8. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe8. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe2. |
| (config-if)#mpls ldp-igp sync isis level-1 | Configure LDP-IGP Synchronization for interface xe8 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be |

| | |
|--|---|
| | <p>acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command.</p> <p>IS-IS: This command is part of ISIS Process.</p> <p>Default: Mandatory configuration. No default option.</p> |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface xe0</code> | Enter interface mode. |
| <code>(config-if)#ip address 70.70.70.1/24</code> | Configure IPv4 address for xe0. |
| <code>(config-if)#ip router isis 1</code> | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe0. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP for IPv4 on xe0. |
| <code>(config-if)#mpls ldp-igp sync isis level-1</code> | <p>Configure LDP-IGP Synchronization for interface xe0 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command.</p> <p>IS-IS: This command is part of ISIS Process.</p> <p>Default: Mandatory configuration. No default option.</p> |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#router isis 1</code> | Create an IS-IS routing instance for area 49 with instance 1 |
| <code>(config-router)#net 49.0001.0000.0000.0002.00</code> | Establish a Network Entity Title for this instance, specifying the area address and the system ID |
| <code>(config-router)#dynamic-hostname</code> | Configure the hostname to be advertised for an ISIS instance |
| <code>(config-router)#is-type level-1</code> | Configure instance as level-1-only routing. |
| <code>(config-router)#metric-style wide</code> | Configure the new style of metric type as wide |
| <code>(config-router)#mpls traffic-eng level-1</code> | Enable MPLS-TE in is-type Level-1. |
| <code>(config-router)#bfd all-interfaces</code> | Enable BFD for ISIS on all interfaces |
| <code>(config-router)#capability cspf</code> | Enable CSPF feature for ISIS instance. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#bfd interval 3 minrx 3 multiplier 3</code> | Configure BFD interval |
| <code>(config)#commit</code> | Commit all the configurations |


P2

| | |
|----------------------------------|---------------------------|
| <code>#configure terminal</code> | Enter configuration mode. |
|----------------------------------|---------------------------|

| | |
|--|--|
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 3.3.3.3/32 secondary | Set the IP address of the loopback interface to 3.3.3.3/32 |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 3.3.3.3 | Set the router ID to IP address 3.3.3.3 |
| (config-router)#transport-address ipv4 3.3.3.3 0 | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. <div>  Note: It is preferable to use the loopback address as the transport address. </div> |
| (config-router)#exit | Exit router mode |
| (config)#interface xe4 | Enter interface mode. |
| (config-if)#ip address 20.20.20.2/24 | Configure IPv4 address for xe4. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe4. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe4. |
| (config-if)#mpls ldp-igp sync isis level-1 | Configure LDP-IGP Synchronization for interface xe4 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command. IS-IS: This command is part of ISIS Process. Default: Mandatory configuration. No default option. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter interface mode. |
| (config-if)#ip address 50.50.50.1/24 | Configure IPv4 address for xe5 |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe5. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe5. |

| | |
|---|--|
| (config-if)#mpls ldp-igp sync isis level-1 | Configure LDP-IGP Synchronization for interface xe5 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command. IS-IS: This command is part of ISIS Process. Default: Mandatory configuration. No default option. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 70.70.70.2/24 | Configure IPv4 address for xe2 |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe0. |
| (config-if)#mpls ldp-igp sync isis level-1 | Configure LDP-IGP Synchronization for interface xe2 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command. IS-IS: This command is part of ISIS Process. Default: Mandatory configuration. No default option. |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 with instance 1 |
| (config-router)#net 49.0001.0000.0000.0003.00 | Establish a Network Entity Title for this instance, specifying the area address and the system ID |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config-router)#is-type level-1 | Configure instance as level-1-only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1. |
| (config-router)#bfd all-interfaces | Enable BFD for ISIS on all interfaces |
| (config-router)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config-if)#exit | Exit interface mode. |
| (config)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval |
| (config)#commit | Commit all the configurations |


P3

| | |
|--|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 4.4.4.4/32 secondary | Set the IP address of the loopback interface to 4.4.4.4/32 |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 4.4.4.4 | Set the router ID to IP address 4.4.4.4 |
| (config-router)#transport-address ipv4 4.4.4.4 0 | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. <div>  Note: It is preferable to use the loopback address as the transport address. </div> |
| (config-router)#exit | Exit router mode |
| (config)#interface xe7 | Enter interface mode. |
| (config-if)#ip address 30.30.30.2/24 | Configure IPv4 address for xe27. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe27. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe27. |
| (config-if)#mpls ldp-igp sync isis level-1 | Configure LDP-IGP Synchronization for interface xe27 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command. IS-IS: This command is part of ISIS Process. Default: Mandatory configuration. No default option. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe4 | Enter interface mode. |
| (config-if)#ip address 60.60.60.1/24 | Configure IPv4 address for xe4. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe4. |

| | |
|---|--|
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe4. |
| (config-if)#mpls ldp-igp sync isis level-1 | Configure LDP-IGP Synchronization for interface xe4 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command. IS-IS: This command is part of ISIS Process. Default: Mandatory configuration. No default option. |
| (config-if)#exit | Exit interface mode |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 with instance 1 |
| (config-router)#net 49.0001.0000.0000.0004.00 | Establish a Network Entity Title for this instance, specifying the area address and the system ID |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config-router)#is-type level-1 | Configure instance as level-1-only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1. |
| (config-router)#bfd all-interfaces | Enable BFD for ISIS on all interfaces |
| (config-router)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config-if)#exit | Exit interface mode. |
| (config)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval |
| (config)#commit | Commit all the configurations |

PE2

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 5.5.5.5/32 secondary | Set the IP address of the loopback interface to 5.5.5.5/32 |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 5.5.5.5 | Set the router ID to IP address 5.5.5.5 |
| (config-router)#transport-address ipv4 5.5.5.5 0 | Configure the transport address for IPv4 (for IPv6 use ipv6) to be used for a TCP session over which LDP will run. |

| | |
|--|---|
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |
| (config)#interface xe9 | Enter interface mode. |
| (config-if)#ip address 40.40.40.2/24 | Configure IPv4 address for xe9. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe9. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe9. |
| (config-if)#mpls ldp-igp sync isis level-1 | <p>Configure LDP-IGP Synchronization for interface xe9 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command.</p> <p>IS-IS: This command is part of ISIS Process.</p> <p>Default: Mandatory configuration. No default option.</p> |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter interface mode. |
| (config-if)#ip address 50.50.50.2/24 | Configure IPv4 address for xe5. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe5. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe5. |
| (config-if)#mpls ldp-igp sync isis level-1 | <p>Configure LDP-IGP Synchronization for interface xe5 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command.</p> <p>IS-IS: This command is part of ISIS Process.</p> <p>Default: Mandatory configuration. No default option.</p> |
| (config-if)#exit | Exit interface mode |

| | |
|---|--|
| (config)#interface xe7 | Enter interface mode. |
| (config-if)#ip address 60.60.60.2/24 | Configure IPv4 address for xe7. |
| (config-if)#ip router isis 1 | Configure IS-IS IPv4 routing on the interface with IS-IS tag instance 1. |
| (config-if)#label-switching | Enable label switching on interface xe7. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe7. |
| (config-if)#mpls ldp-igp sync isis level-1 | Configure LDP-IGP Synchronization for interface xe7 belonging to an IS-IS process with corresponding IS-IS level. The values level-1 level-2-only level-1-2 identify the IS-IS level instance. The interface can be acting on any level, but the sync is applicable only when it matches with the level given in IGP sync command. IS-IS: This command is part of ISIS Process. Default: Mandatory configuration. No default option. |
| (config-if)#exit | Exit interface mode |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 with instance 1 |
| (config-router)#net 49.0001.0000.0000.0005.00 | Establish a Network Entity Title for this instance, specifying the area address and the system ID |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config-router)#is-type level-1 | Configure instance as level-1-only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1. |
| (config-router)#bfd all-interfaces | Enable BFD for ISIS on all interfaces |
| (config-router)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config-if)#exit | Exit interface mode. |
| (config)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval |
| (config)#commit | Commit all the configurations |

Validation

PE1

```
PE1#show clns neighbors
```

```
Total number of L1 adjacencies: 3
```

```
Total number of L2 adjacencies: 0
```

```
Total number of adjacencies: 3
```

```
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P1 | ce46/1 | b86a.97f5.ead2 | Up | 8 | L1 | IS-IS |
| P3 | ce46/3 | e8c5.7a8c.f2c8 | Up | 9 | L1 | IS-IS |
| P2 | ce46/4 | e49d.731b.ecc5 | Up | 6 | L1 | IS-IS |

```
PE1#show clns neighbors detail
```

```

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1: VRF : default
System Id      Interface  SNPA              State  Holdtime  Type Protocol
P1             ce46/1    b86a.97f5.ead2    Up     8          L1    IS-IS
  L1 Adjacency ID: 1
  L2 Adjacency ID: 2
  Uptime: 00:17:48
  Area Address(es): 49.0001
  IP Address(es): 10.10.10.2
  Level-1 Protocols Supported: IPv4
  IPV4 Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

P3             ce46/3    e8c5.7a8c.f2c8    Up     8          L1    IS-IS
  L1 Adjacency ID: 1
  L2 Adjacency ID: 2
  Uptime: 00:45:52
  Area Address(es): 49.0001
  IP Address(es): 30.30.30.2

...skipping 1 line
  IPV4 Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

P2             ce46/4    e49d.731b.ecc5    Up     6          L1    IS-IS
  L1 Adjacency ID: 1
  L2 Adjacency ID: 2
  Uptime: 00:47:07
  Area Address(es): 49.0001
  IP Address(es): 20.20.20.2
  Level-1 Protocols Supported: IPv4
  IPV4 Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

PE1#how isis database
^
% Invalid input detected at '^' marker.

PE1#show isis database verbose
Tag 1: VRF : default
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00      * 0x00000009 0xA301        1033          0/0/0
  Area Address: 49.0001
  NLPID:        0xCC
  Hostname:     PE1
  IP Address:   1.1.1.1
  Router ID:    1.1.1.1
  Metric:       15      IS-Extended P2.02
    IPv4 Interface Address: 20.20.20.1
    Neighbor IP Address: 20.20.20.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec

```

```

TE-Default Metric: 15
Metric: 60 IS-Extended P3.02
IPv4 Interface Address: 30.30.30.1
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec

...skipping 1 line
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 60
Metric: 20 IS-Extended P1.04
IPv4 Interface Address: 10.10.10.1
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 20
Metric: 10 IP-Extended 1.1.1.1/32
Prefix Attribute Flags[0]: ELC Set
Metric: 60 IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 15 IP-Extended 20.20.20.0/24
Most commands optionally preceded by integer argument k. Defaults in brackets.
Star (*) indicates argument becomes new default.
-----
<space>          Display next k lines of text [current screen size]
z                Display next k lines of text [current screen size]*
<return>        Display next k lines of text [1]*
d or ctrl-D      Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s                Skip forward k lines of text [1]
f                Skip forward k screenfuls of text [1]
b or ctrl-B      Skip backwards k screenfuls of text [1]
'                Go to place where previous search started
=                Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n                Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v                Start up '/usr/bin/vi' at current line
ctrl-L           Redraw screen
:n               Go to kth next file [1]
:p               Go to kth previous file [1]
:f               Display current file name and line number
.               Repeat previous command
-----
Prefix Attribute Flags[0]: ELC Set
Metric: 20 IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
Pl.00-00          0x00000007 0xFB9F 1036 0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P1
IP Address: 2.2.2.2

```

```

Router ID:      2.2.2.2
Metric: 10      IS-Extended P2.04
IPv4 Interface Address: 70.70.70.1
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P1.02
IPv4 Interface Address: 40.40.40.1
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P1.04
IPv4 Interface Address: 10.10.10.2
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 2.2.2.2/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
P1.02-00      0x00000003  0xC103      391      0/0/0
Metric: 0      IS-Extended P1.00
Metric: 0      IS-Extended PE2.00
P1.04-00      0x00000002  0x6562      1032     0/0/0
Metric: 0      IS-Extended P1.00
Metric: 0      IS-Extended PE1.00
P2.00-00      0x00000006  0xF247      391      0/0/0

...skipping 1 line
NLPID:      0xCC
Hostname:    P2
IP Address:  3.3.3.3
Router ID:   3.3.3.3
Metric: 10   IS-Extended P2.02

```

```

IPv4 Interface Address: 20.20.20.2
Neighbor IP Address: 20.20.20.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P2.04
IPv4 Interface Address: 70.70.70.2
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.04
IPv4 Interface Address: 80.80.80.2
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec

...skipping 23 lines
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P2.02-00      0x00000004  0x7055      1074      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended PE1.00
P2.03-00      0x00000003  0xBB06      391      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended PE2.00
P2.04-00      0x00000004  0x764C      1074      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended P1.00
P3.00-00      0x00000005  0x74F0      391      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P3
IP Address: 4.4.4.4
Router ID: 4.4.4.4
Metric: 10      IS-Extended P3.02
IPv4 Interface Address: 30.30.30.2
Neighbor IP Address: 30.30.30.2

...skipping 23 lines

```



```

    Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.03
IPv4 Interface Address: 60.60.60.1
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 4.4.4.4/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P3.02-00      0x00000004  0x7152      1149      0/0/0
Metric: 0      IS-Extended P3.00
Metric: 0      IS-Extended PE1.00
P3.03-00      0x00000003  0xBC03      391      0/0/0
Metric: 0      IS-Extended P3.00
Metric: 0      IS-Extended PE2.00
P3.04-00      0x00000004  0x8B34      1149      0/0/0
Metric: 0      IS-Extended P3.00
Metric: 0      IS-Extended P2.00
PE2.00-00      0x00000004  0x8799      395      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: PE2
IP Address: 5.5.5.5
Router ID: 5.5.5.5
Metric: 10      IS-Extended P2.03
IPv4 Interface Address: 50.50.50.2
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.03
IPv4 Interface Address: 60.60.60.2
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec

```

```

    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec

...skipping 1 line
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended Pl.02
IPv4 Interface Address: 40.40.40.2
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 5.5.5.5/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set

PE1#p routes
%% Bad IP address or hostname
PE1#show ldp fec
LSR codes      : E/N - LSR is egress/non-egress for this FEC,
                  L - LSR received a label for this FEC,
                  P - Primary route, B - LFA Backup route,
                  R - Remote LFA Backup route,
                  > - LSR will use this route for the FEC
FEC             Code   Session      Out Label   ELC   Nexthop Addr
1.1.1.1/32      NL     2.2.2.2     24328       No    connected
                NL     4.4.4.4     24327       No    connected
                NL     3.3.3.3     24326       No    connected
                E >    non-existent none         No    connected
                NL     5.5.5.5     24327       No    connected
2.2.2.2/32      NLP>   2.2.2.2     impl-null   No    10.10.10.2
                NL     4.4.4.4     24324       No    no nexthop
                NL     3.3.3.3     24323       No    no nexthop
                NL     5.5.5.5     24320       No    no nexthop
3.3.3.3/32      NL     2.2.2.2     24320       No    no nexthop
                NL     4.4.4.4     24320       No    no nexthop
                NLP>   3.3.3.3     impl-null   No    20.20.20.2
                NL     5.5.5.5     24321       No    no nexthop
4.4.4.4/32      NL     2.2.2.2     24321       No    no nexthop
                NLP>   3.3.3.3     24320       No    20.20.20.2
                NL     4.4.4.4     impl-null   No    no nexthop
                NL     5.5.5.5     24322       No    no nexthop
5.5.5.5/32      NL     2.2.2.2     24322       No    no nexthop
                NL     4.4.4.4     24325       No    no nexthop
                NLP>   3.3.3.3     24324       No    20.20.20.2
                NL     5.5.5.5     24328       No    no nexthop
10.10.10.0/24   NL     2.2.2.2     impl-null   No    connected
                NL     4.4.4.4     24328       No    connected
                NL     3.3.3.3     24327       No    connected
                E >    non-existent none         No    connected
20.20.20.0/24   NL     2.2.2.2     24323       No    connected
                NL     3.3.3.3     impl-null   No    connected
                NL     4.4.4.4     24323       No    connected

...skipping 1 line

```

```

E >      non-existent      none      No      connected
30.30.30.0/24 NL      2.2.2.2      24324      No      connected
NL      4.4.4.4      impl-null      No      connected
NL      3.3.3.3      24322      No      connected
NL      5.5.5.5      24324      No      connected
E >      non-existent      none      No      connected
40.40.40.0/24 NLP>    2.2.2.2      impl-null      No      10.10.10.2
NL      4.4.4.4      24326      No      no nexthop
NL      3.3.3.3      24325      No      no nexthop
NL      5.5.5.5      24329      No      no nexthop
50.50.50.0/24 NL      2.2.2.2      24325      No      no nexthop
NL      4.4.4.4      24321      No      no nexthop
NLP>    3.3.3.3      impl-null      No      20.20.20.2
NL      5.5.5.5      24330      No      no nexthop
60.60.60.0/24 NL      2.2.2.2      24326      No      no nexthop
NLP>    3.3.3.3      24321      No      20.20.20.2
NL      4.4.4.4      impl-null      No      no nexthop
NL      5.5.5.5      24331      No      no nexthop
70.70.70.0/24 NL      2.2.2.2      impl-null      No      no nexthop
NL      4.4.4.4      24322      No      no nexthop
NLP>    3.3.3.3      impl-null      No      20.20.20.2
NL      5.5.5.5      24325      No      no nexthop
80.80.80.0/24 NL      2.2.2.2      24327      No      no nexthop

```

Most commands optionally preceded by integer argument k. Defaults in brackets.
 Star (*) indicates argument becomes new default.

```

-----
<space>      Display next k lines of text [current screen size]
z            Display next k lines of text [current screen size]*
<return>    Display next k lines of text [1]*
d or ctrl-D  Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s            Skip forward k lines of text [1]
f            Skip forward k screenfuls of text [1]
b or ctrl-B  Skip backwards k screenfuls of text [1]
'            Go to place where previous search started
=            Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n            Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v            Start up '/usr/bin/vi' at current line
ctrl-L       Redraw screen
:n           Go to kth next file [1]
:p           Go to kth previous file [1]
:f           Display current file name and line number
.            Repeat previous command
-----

```

...skipping 1 line

```

NL      4.4.4.4      impl-null      No      no nexthop
NL      5.5.5.5      24326      No      no nexthop

```

PE1#

PE1#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|------------|---------|----------|------------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 2.2.2.2/32 | 1 | 17 | - | - | - | - |
| | - | N/A | 01:30:30 | - | - | - | - |
| | | | 15 | - | - | - | - |
| | Yes 3 | ce46/1 | No | 10.10.10.2 | - | - | - |
| L> | 3.3.3.3/32 | 4 | 8 | - | - | - | - |
| | - | N/A | 01:30:30 | - | - | - | - |
| | | | 16 | - | - | - | - |
| | Yes 3 | ce46/4 | No | 20.20.20.2 | - | - | - |

```

L> 4.4.4.4/32      5      9      -      -      -      -      -
    -      N/A      01:30:30
                        62
    Yes 24320      ce46/4      No      20.20.20.2      -      -
L> 5.5.5.5/32      2      11      -      -      -      -      -
    -      N/A      01:30:30
                        4
    Yes 24324      ce46/4      No      20.20.20.2      -      -
L> 40.40.40.0/24   3      17      -      -      -      -      -
    -      N/A      01:30:30
                        15
    Yes 3      ce46/1      No      10.10.10.2      -      -
L> 50.50.50.0/24   8      8      -      -      -      -      -
    -      N/A      01:30:30
                        16
    Yes 3      ce46/4      No      20.20.20.2      -      -
L> 60.60.60.0/24   6      14      -      -      -      -      -
    -      N/A      01:30:30
                        72
    Yes 24321      ce46/4      No      20.20.20.2      -      -
L> 70.70.70.0/24   9      8      -      -      -      -      -
    -      N/A      01:30:30
                        16
    Yes 3      ce46/4      No      20.20.20.2      -      -
L> 80.80.80.0/24   7      8      -      -      -      -      -
    -      N/A      01:30:30
                        16
    Yes 3      ce46/4      No      20.20.20.2      -      -
PE1#s forwarding-table 5.5.5.5/32
    ^
% Invalid input detected at '^' marker.

PE1#show mpls ftn-table
Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:30:31, UpTime: 01:30:31, LastUpdate: 00:17:46
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 15 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 15, owner: N/A, Stale: NO, refcount: 4, out intf: ce46/1, out label: 3
Nexthop addr: 10.10.10.2      cross connect ix: 13, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:30:31, UpTime: 01:30:31, LastUpdate: 00:27:53
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2      cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:30:31, UpTime: 01:30:31, LastUpdate: 00:27:53
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 62 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 62, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24320
Nexthop addr: 20.20.20.2      cross connect ix: 11, op code: Push

```

```
Primary FTN entry with FEC: 5.5.5.5/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:30:31, UpTime: 01:30:31, LastUpdate: 00:27:53
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 4, owner: LDP, Stale: NO, refcount: 4, out intf: ce46/4, out label: 24324
  Nexthop addr: 20.20.20.2      cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 40.40.40.0/24, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:30:31, UpTime: 01:30:31, LastUpdate: 00:17:46
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 15 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 15, owner: N/A, Stale: NO, refcount: 4, out intf: ce46/1, out label: 3
  Nexthop addr: 10.10.10.2      cross connect ix: 13, op code: Push

Primary FTN entry with FEC: 50.50.50.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:30:31, UpTime: 01:30:31, LastUpdate: 00:27:53
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
  Nexthop addr: 20.20.20.2      cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 60.60.60.0/24, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:30:31, UpTime: 01:30:31, LastUpdate: 00:27:53
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 72 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 72, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24321
  Nexthop addr: 20.20.20.2      cross connect ix: 12, op code: Push

Primary FTN entry with FEC: 70.70.70.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:30:31, UpTime: 01:30:31, LastUpdate: 00:27:53
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
  Nexthop addr: 20.20.20.2      cross connect ix: 6, op code: Push

...skipping 1 line
  CreateTime: 01:30:31, UpTime: 01:30:31, LastUpdate: 00:27:53
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
 Nexthop addr: 20.20.20.2 cross connect ix: 6, op code: Push

PE1#show clns neighbors

^

% Invalid input detected at '^' marker.

PE1#show clns neighbors detail

Total number of L1 adjacencies: 3

Total number of L2 adjacencies: 0

Total number of adjacencies: 3

Tag 1: VRF : default

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P1 | ce46/1 | b86a.97f5.ead2 | Up | 8 | L1 | IS-IS |

L1 Adjacency ID: 1

L2 Adjacency ID: 2

Uptime: 00:17:51

Area Address(es): 49.0001

IP Address(es): 10.10.10.2

Level-1 Protocols Supported: IPv4

IPv4 Bidirectional Forwarding Detection is enabled

Adjacency advertisement: Advertise

| | | | | | | |
|----|--------|----------------|----|---|----|-------|
| P3 | ce46/3 | e8c5.7a8c.f2c8 | Up | 9 | L1 | IS-IS |
|----|--------|----------------|----|---|----|-------|

L1 Adjacency ID: 1

L2 Adjacency ID: 2

Uptime: 00:45:55

Area Address(es): 49.0001

IP Address(es): 30.30.30.2

...skipping 1 line

IPv4 Bidirectional Forwarding Detection is enabled

Adjacency advertisement: Advertise

| | | | | | | |
|----|--------|----------------|----|---|----|-------|
| P2 | ce46/4 | e49d.731b.ecc5 | Up | 6 | L1 | IS-IS |
|----|--------|----------------|----|---|----|-------|

L1 Adjacency ID: 1

L2 Adjacency ID: 2

Uptime: 00:47:10

Area Address(es): 49.0001

IP Address(es): 20.20.20.2

Level-1 Protocols Supported: IPv4

IPv4 Bidirectional Forwarding Detection is enabled

Adjacency advertisement: Advertise

PE1#s database

^

% Invalid input detected at '^' marker.

PE1#show isis database verbose

Tag 1: VRF : default

IS-IS Level-1 Link State Database:

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|-----------|--------------|--------------|--------------|----------|
| PE1.00-00 | * 0x00000009 | 0xA301 | 1030 | 0/0/0 |

Area Address: 49.0001

NLPID: 0xCC

Hostname: PE1

IP Address: 1.1.1.1

Router ID: 1.1.1.1

Metric: 15 IS-Extended P2.02

IPv4 Interface Address: 20.20.20.1

Neighbor IP Address: 20.20.20.2

Maximum Link Bandwidth : 10000000.00 kbits/sec

Reservable Bandwidth : 10000000.00 kbits/sec

```

Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 15
Metric: 60          IS-Extended P3.02
IPv4 Interface Address: 30.30.30.1
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 60
Metric: 20          IS-Extended P1.04
IPv4 Interface Address: 10.10.10.1
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 20
Metric: 10          IP-Extended 1.1.1.1/32
Prefix Attribute Flags[0]: ELC Set
Metric: 60          IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 15          IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 20          IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
P1.00-00          0x00000007  0xFB9F          1033          0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P1
IP Address: 2.2.2.2
Router ID: 2.2.2.2
Metric: 10          IS-Extended P2.04
IPv4 Interface Address: 70.70.70.1
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec

```

```

TE-Default Metric: 10
Metric: 10      IS-Extended P1.02
IPv4 Interface Address: 40.40.40.1
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P1.04
IPv4 Interface Address: 10.10.10.2
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 2.2.2.2/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
P1.02-00      0x00000003  0xC103      387      0/0/0
Metric: 0      IS-Extended P1.00
Metric: 0      IS-Extended PE2.00
P1.04-00      0x00000002  0x6562      1029     0/0/0
Metric: 0      IS-Extended P1.00
Metric: 0      IS-Extended PE1.00
P2.00-00      0x00000006  0xF247      388      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P2
IP Address: 3.3.3.3
Router ID: 3.3.3.3
Metric: 10      IS-Extended P2.02
IPv4 Interface Address: 20.20.20.2
Neighbor IP Address: 20.20.20.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P2.04
IPv4 Interface Address: 70.70.70.2

```



```

Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10 IS-Extended P3.04
IPv4 Interface Address: 80.80.80.2
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10 IS-Extended P2.03
IPv4 Interface Address: 50.50.50.1
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10 IP-Extended 3.3.3.3/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P2.02-00 0x00000004 0x7055 1070 0/0/0
Metric: 0 IS-Extended P2.00
Metric: 0 IS-Extended PE1.00
P2.03-00 0x00000003 0xBB06 388 0/0/0
Metric: 0 IS-Extended P2.00
Metric: 0 IS-Extended PE2.00
P2.04-00 0x00000004 0x764C 1070 0/0/0
Metric: 0 IS-Extended P2.00
Metric: 0 IS-Extended P1.00
P3.00-00 0x00000005 0x74F0 388 0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P3
IP Address: 4.4.4.4
Router ID: 4.4.4.4

```

```

Metric: 10          IS-Extended P3.02
IPv4 Interface Address: 30.30.30.2
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.04
IPv4 Interface Address: 80.80.80.1
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.03
IPv4 Interface Address: 60.60.60.1
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 4.4.4.4/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P3.02-00            0x00000004  0x7152      1146      0/0/0
Metric: 0           IS-Extended P3.00
Metric: 0           IS-Extended PE1.00
P3.03-00            0x00000003  0xBC03      388      0/0/0
Metric: 0           IS-Extended P3.00
Metric: 0           IS-Extended PE2.00
P3.04-00            0x00000004  0x8B34      1146      0/0/0
Metric: 0           IS-Extended P3.00
Metric: 0           IS-Extended P2.00
PE2.00-00           0x00000004  0x8799      391      0/0/0
Area Address: 49.0001
NLPID:              0xCC
Hostname:           PE2
IP Address:         5.5.5.5
Router ID:          5.5.5.5

```

```

Metric: 10          IS-Extended P2.03
IPv4 Interface Address: 50.50.50.2
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.03
IPv4 Interface Address: 60.60.60.2
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P1.02
IPv4 Interface Address: 40.40.40.2
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 5.5.5.5/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set

PE1#ping 5.5.5.5
Press CTRL+C to exit
PING 5.5.5.5 (5.5.5.5) 100(128) bytes of data.
108 bytes from 5.5.5.5: icmp_seq=1 ttl=64 time=0.602 ms
108 bytes from 5.5.5.5: icmp_seq=2 ttl=64 time=0.508 ms
108 bytes from 5.5.5.5: icmp_seq=3 ttl=64 time=0.624 ms
108 bytes from 5.5.5.5: icmp_seq=4 ttl=64 time=0.384 ms
108 bytes from 5.5.5.5: icmp_seq=5 ttl=64 time=0.484 ms

--- 5.5.5.5 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4086ms
rtt min/avg/max/mdev = 0.384/0.520/0.624/0.086 ms

```

PE2

```
PE2#show clns neighbors
```

```
Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
```

```
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P2 | xe5 | e49d.731b.ecc6 | Up | 7 | L1 | IS-IS |
| P3 | xe7 | e8c5.7a8c.f2c5 | Up | 7 | L1 | IS-IS |
| P1 | xe9 | b86a.97f5.eace | Up | 7 | L1 | IS-IS |

```
PE2#show clns neighbors detail
```

```
Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
```

```
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P2 | xe5 | e49d.731b.ecc6 | Up | 7 | L1 | IS-IS |

```
L1 Adjacency ID: 1
```

```
L2 Adjacency ID: 2
```

```
Uptime: 00:45:14
```

```
Area Address(es): 49.0001
```

```
IP Address(es): 50.50.50.1
```

```
Level-1 Protocols Supported: IPv4
```

```
IPV4 Bidirectional Forwarding Detection is enabled
```

```
Adjacency advertisement: Advertise
```

| | | | | | | |
|----|-----|----------------|----|---|----|-------|
| P3 | xe7 | e8c5.7a8c.f2c5 | Up | 7 | L1 | IS-IS |
|----|-----|----------------|----|---|----|-------|

```
L1 Adjacency ID: 1
```

```
L2 Adjacency ID: 2
```

```
Uptime: 00:45:14
```

```
Area Address(es): 49.0001
```

```
IP Address(es): 60.60.60.1
```

```
...skipping 1 line
```

```
IPV4 Bidirectional Forwarding Detection is enabled
```

```
Adjacency advertisement: Advertise
```

| | | | | | | |
|----|-----|----------------|----|---|----|-------|
| P1 | xe9 | b86a.97f5.eace | Up | 7 | L1 | IS-IS |
|----|-----|----------------|----|---|----|-------|

```
L1 Adjacency ID: 1
```

```
L2 Adjacency ID: 2
```

```
Uptime: 00:45:14
```

```
Area Address(es): 49.0001
```

```
IP Address(es): 40.40.40.1
```

```
Level-1 Protocols Supported: IPv4
```

```
IPV4 Bidirectional Forwarding Detection is enabled
```

```
Adjacency advertisement: Advertise
```

```
PE2#s database
```

```
^
```

```
% Invalid input detected at '^' marker.
```

```
PE2#show isis database verbose
```

```
Tag 1: VRF : default
```

```
IS-IS Level-1 Link State Database:
```

| LSPID | LSP Seq Num | LSP Checksum | LSP Holdtime | ATT/P/OL |
|-----------|-------------|--------------|--------------|----------|
| PE1.00-00 | 0x00000009 | 0xA301 | 929 | 0/0/0 |

```
Area Address: 49.0001
NLPID: 0xCC
Hostname: PE1
IP Address: 1.1.1.1
Router ID: 1.1.1.1
Metric: 15 IS-Extended P2.02
```

```

IPv4 Interface Address: 20.20.20.1
Neighbor IP Address: 20.20.20.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 15
Metric: 60      IS-Extended P3.02
IPv4 Interface Address: 30.30.30.1
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
Most commands optionally preceded by integer argument k. Defaults in brackets.
Star (*) indicates argument becomes new default.
-----
<space>          Display next k lines of text [current screen size]
z                Display next k lines of text [current screen size]*
<return>         Display next k lines of text [1]*
d or ctrl-D      Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s                Skip forward k lines of text [1]
f                Skip forward k screenfuls of text [1]
b or ctrl-B      Skip backwards k screenfuls of text [1]
'                Go to place where previous search started
=                Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n                Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v                Start up '/usr/bin/vi' at current line
ctrl-L           Redraw screen
:n               Go to kth next file [1]
:p               Go to kth previous file [1]
:f               Display current file name and line number
.               Repeat previous command
-----
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 60
Metric: 20      IS-Extended P1.04
IPv4 Interface Address: 10.10.10.1
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 20
Metric: 10      IP-Extended 1.1.1.1/32

```

```

Prefix Attribute Flags[0]: ELC Set
Metric: 60 IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 15 IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 20 IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
P1.00-00 0x00000007 0xFB9F 933 0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P1
IP Address: 2.2.2.2
Router ID: 2.2.2.2
Metric: 10 IS-Extended P2.04
IPv4 Interface Address: 70.70.70.1
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec

...skipping 1 line
Metric: 10 IS-Extended P1.02
IPv4 Interface Address: 40.40.40.1
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10 IS-Extended P1.04
IPv4 Interface Address: 10.10.10.2
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
Most commands optionally preceded by integer argument k. Defaults in brackets.
Star (*) indicates argument becomes new default.
-----
<space>      Display next k lines of text [current screen size]
z            Display next k lines of text [current screen size]*
<return>    Display next k lines of text [1]*
d or ctrl-D  Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s            Skip forward k lines of text [1]
f            Skip forward k screenfuls of text [1]
b or ctrl-B  Skip backwards k screenfuls of text [1]
'            Go to place where previous search started
=            Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n            Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell

```

```

v          Start up '/usr/bin/vi' at current line
ctrl-L     Redraw screen
:n         Go to kth next file [1]
:p         Go to kth previous file [1]
:f         Display current file name and line number
.          Repeat previous command
-----

    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric:    10      IP-Extended 2.2.2.2/32
Prefix Attribute Flags[0]: ELC Set
Metric:    10      IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric:    10      IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set
Metric:    10      IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
Pl.02-00    0x00000004  0xBF04      1187      0/0/0
Metric:     0      IS-Extended Pl.00
Metric:     0      IS-Extended PE2.00
Pl.04-00    0x00000002  0x6562      929      0/0/0
Metric:     0      IS-Extended Pl.00
Metric:     0      IS-Extended PE1.00
P2.00-00    0x00000007  0xF048      1188      0/0/0
Area Address: 49.0001
Most commands optionally preceded by integer argument k.  Defaults in brackets.
Star (*) indicates argument becomes new default.
-----

<space>     Display next k lines of text [current screen size]
z           Display next k lines of text [current screen size]*
<return>    Display next k lines of text [1]*
d or ctrl-D Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s           Skip forward k lines of text [1]
f           Skip forward k screenfuls of text [1]
b or ctrl-B Skip backwards k screenfuls of text [1]
'           Go to place where previous search started
=           Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n           Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v           Start up '/usr/bin/vi' at current line
ctrl-L     Redraw screen
:n         Go to kth next file [1]
:p         Go to kth previous file [1]
:f         Display current file name and line number
.          Repeat previous command
-----

...skipping 1 line
Hostname:    P2
IP Address:  3.3.3.3
Router ID:   3.3.3.3
Metric:     10      IS-Extended P2.02
IPv4 Interface Address: 20.20.20.2
Neighbor IP Address: 20.20.20.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec

```

```

    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P2.04
IPv4 Interface Address: 70.70.70.2
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.04
IPv4 Interface Address: 80.80.80.2
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec

...skipping 1 line
TE-Default Metric: 10
Metric: 10          IS-Extended P2.03
IPv4 Interface Address: 50.50.50.1
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 3.3.3.3/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 70.70.70.0/24
Most commands optionally preceded by integer argument k. Defaults in brackets.
Star (*) indicates argument becomes new default.
-----
<space>          Display next k lines of text [current screen size]
z                Display next k lines of text [current screen size]*
<return>         Display next k lines of text [1]*
d or ctrl-D      Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s                Skip forward k lines of text [1]
f                Skip forward k screenfuls of text [1]

```



```

b or ctrl-B      Skip backwards k screenfuls of text [1]
'               Go to place where previous search started
=               Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n               Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v               Start up '/usr/bin/vi' at current line
ctrl-L          Redraw screen
:n              Go to kth next file [1]
:p              Go to kth previous file [1]
:f              Display current file name and line number
.               Repeat previous command
-----
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P2.02-00        0x00000004 0x7055      970      0/0/0
Metric: 0        IS-Extended P2.00
Metric: 0        IS-Extended PE1.00
P2.03-00        0x00000004 0xB907      1187     0/0/0
Metric: 0        IS-Extended P2.00
Metric: 0        IS-Extended PE2.00
P2.04-00        0x00000004 0x764C      970      0/0/0
Metric: 0        IS-Extended P2.00
Metric: 0        IS-Extended P1.00
P3.00-00        0x00000006 0x72F1      1188     0/0/0
Area Address: 49.0001
NLPID:          0xCC
Hostname:       P3
IP Address:     4.4.4.4
Router ID:      4.4.4.4
Metric: 10      IS-Extended P3.02
IPv4 Interface Address: 30.30.30.2
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Most commands optionally preceded by integer argument k. Defaults in brackets.
Star (*) indicates argument becomes new default.
-----
<space>         Display next k lines of text [current screen size]
z               Display next k lines of text [current screen size]*
<return>        Display next k lines of text [1]*
d or ctrl-D     Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s               Skip forward k lines of text [1]
f               Skip forward k screenfuls of text [1]
b or ctrl-B     Skip backwards k screenfuls of text [1]
'               Go to place where previous search started
=               Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n               Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v               Start up '/usr/bin/vi' at current line
ctrl-L          Redraw screen
:n              Go to kth next file [1]
:p              Go to kth previous file [1]
:f              Display current file name and line number
.               Repeat previous command
-----
...skipping 23 lines
Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.03
IPv4 Interface Address: 60.60.60.1
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec

```

```

Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 4.4.4.4/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P3.02-00      0x00000004 0x7152      1046      0/0/0
Metric: 0      IS-Extended P3.00
Metric: 0      IS-Extended PE1.00
P3.03-00      0x00000004 0xBA04      1188      0/0/0
Metric: 0      IS-Extended P3.00
Metric: 0      IS-Extended PE2.00
P3.04-00      0x00000004 0x8B34      1046      0/0/0
Metric: 0      IS-Extended P3.00
Metric: 0      IS-Extended P2.00
PE2.00-00      * 0x00000005 0x859A      1192      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: PE2
IP Address: 5.5.5.5
Router ID: 5.5.5.5
Metric: 10      IS-Extended P2.03
IPv4 Interface Address: 50.50.50.2
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:

...skipping 1 line
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.03
IPv4 Interface Address: 60.60.60.2
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P1.02
IPv4 Interface Address: 40.40.40.2
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec

```

```

Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 5.5.5.5/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set
Most commands optionally preceded by integer argument k.  Defaults in brackets.
Star (*) indicates argument becomes new default.
-----
<space>          Display next k lines of text [current screen size]
z                Display next k lines of text [current screen size]*
<return>         Display next k lines of text [1]*
d or ctrl-D      Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s                Skip forward k lines of text [1]
f                Skip forward k screenfuls of text [1]
b or ctrl-B      Skip backwards k screenfuls of text [1]
'                Go to place where previous search started
=                Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n                Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v                Start up '/usr/bin/vi' at current line
ctrl-L           Redraw screen
:n               Go to kth next file [1]
:p               Go to kth previous file [1]
:f               Display current file name and line number
.                Repeat previous command
-----

...skipping 23 lines
PE2#rwarding-table
^
% Invalid input detected at '^' marker.

PE2#show mpls forwarding-table 5.5.5.5/32
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code   FEC           FTN-ID   Nhlfe-ID  Tunnel-ID   Pri   Out-Label   Out-
Intf   ELC           Nexthop   Algo-Num   UpTime
PE2#show mpls ftn-table
Primary FTN entry with FEC: 1.1.1.1/32, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:32:39, UpTime: 01:32:39, LastUpdate: 00:29:35
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 26 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 26, owner: LDP, Stale: NO, refcount: 2, out intf: xe5, out label: 24326

```

Nexthop addr: 50.50.50.1 cross connect ix: 4, op code: Push

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 18 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 18, owner: LDP, Stale: NO, refcount: 4, out intf: xe7, out label: 24327

Nexthop addr: 60.60.60.1 cross connect ix: 4, op code: Push

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 20 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 20, owner: LDP, Stale: NO, refcount: 2, out intf: xe9, out label: 24328

Nexthop addr: 40.40.40.1 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:33:38, UpTime: 01:33:38, LastUpdate: 00:29:35

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 6, out intf: xe9, out label: 3

Nexthop addr: 40.40.40.1 cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 3, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:33:34, UpTime: 01:33:34, LastUpdate: 00:29:35

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 27 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 27, owner: N/A, Stale: NO, refcount: 11, out intf: xe5, out label: 3

Nexthop addr: 50.50.50.1 cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 6, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:32:39, UpTime: 01:32:39, LastUpdate: 00:29:35

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 22 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 22, owner: N/A, Stale: NO, refcount: 6, out intf: xe7, out label: 3

Nexthop addr: 60.60.60.1 cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 10.10.10.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 00:20:40, UpTime: 00:20:40, LastUpdate: 00:19:33

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 6, out intf: xe9, out label: 3

Nexthop addr: 40.40.40.1 cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 20.20.20.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:33:34, UpTime: 01:33:34, LastUpdate: 00:29:35

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 27 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

```

    Out-segment with ix: 27, owner: N/A, Stale: NO, refcount: 11, out intf: xe5, out label: 3
    Nexthop addr: 50.50.50.1          cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 30.30.30.0/24, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:32:39, UpTime: 01:32:39, LastUpdate: 00:29:35
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 22 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 22, owner: N/A, Stale: NO, refcount: 6, out intf: xe7, out label: 3
  Nexthop addr: 60.60.60.1          cross connect ix: 6, op code: Push

...skipping 1 line
  CreateTime: 01:33:38, UpTime: 01:33:38, LastUpdate: 00:29:35
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 6, out intf: xe9, out label: 3
  Nexthop addr: 40.40.40.1          cross connect ix: 2, op code: Push

  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 27 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 27, owner: N/A, Stale: NO, refcount: 11, out intf: xe5, out label: 3
  Nexthop addr: 50.50.50.1          cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 80.80.80.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:32:39, UpTime: 01:32:39, LastUpdate: 00:29:35
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 22 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 22, owner: N/A, Stale: NO, refcount: 6, out intf: xe7, out label: 3
  Nexthop addr: 60.60.60.1          cross connect ix: 6, op code: Push

...skipping 1 line
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 27, owner: N/A, Stale: NO, refcount: 11, out intf: xe5, out label: 3
  Nexthop addr: 50.50.50.1          cross connect ix: 6, op code: Push

PE2# neighbors
      ^
% Invalid input detected at '^' marker.

PE2#show clns neighbors detail

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1: VRF : default
System Id      Interface  SNPA              State  Holdtime  Type Protocol
P2             xe5         e49d.731b.ecc6    Up     7         L1    IS-IS
  L1 Adjacency ID: 1
  L2 Adjacency ID: 2
  Uptime: 00:45:17
  Area Address(es): 49.0001
  IP Address(es): 50.50.50.1

```

```

Level-1 Protocols Supported: IPv4
IPv4 Bidirectional Forwarding Detection is enabled
Adjacency advertisement: Advertise

P3          xe7          e8c5.7a8c.f2c5      Up      7          L1    IS-IS
L1 Adjacency ID: 1
L2 Adjacency ID: 2
Uptime: 00:45:17
Area Address(es): 49.0001
IP Address(es): 60.60.60.1

...skipping 1 line
IPv4 Bidirectional Forwarding Detection is enabled
Adjacency advertisement: Advertise

P1          xe9          b86a.97f5.eace      Up      7          L1    IS-IS
L1 Adjacency ID: 1
L2 Adjacency ID: 2
Uptime: 00:45:17
Area Address(es): 49.0001
IP Address(es): 40.40.40.1
Level-1 Protocols Supported: IPv4
IPv4 Bidirectional Forwarding Detection is enabled
Adjacency advertisement: Advertise

PE2#s database
^
% Invalid input detected at '^' marker.

PE2#show isis database verbose
Tag 1: VRF : default
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00      0x00000009  0xA301        926           0/0/0
Area Address: 49.0001
NLPID:         0xCC
Hostname:      PE1
IP Address:    1.1.1.1
Router ID:     1.1.1.1
Metric: 15      IS-Extended P2.02
IPv4 Interface Address: 20.20.20.1
Neighbor IP Address: 20.20.20.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 15
Metric: 60      IS-Extended P3.02
IPv4 Interface Address: 30.30.30.1
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec

```

```

    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 60
Metric: 20      IS-Extended P1.04
IPv4 Interface Address: 10.10.10.1
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 20
Metric: 10      IP-Extended 1.1.1.1/32
Prefix Attribute Flags[0]: ELC Set
Metric: 60      IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 15      IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 20      IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
P1.00-00      0x00000007  0xFB9F  930  0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P1
IP Address: 2.2.2.2
Router ID: 2.2.2.2
Metric: 10      IS-Extended P2.04
IPv4 Interface Address: 70.70.70.1
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P1.02
IPv4 Interface Address: 40.40.40.1
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P1.04
IPv4 Interface Address: 10.10.10.2
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec

```

```

Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 2.2.2.2/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
Pl.02-00      0x00000004  0xBF04      1184      0/0/0
Metric: 0      IS-Extended P1.00
Metric: 0      IS-Extended PE2.00
Pl.04-00      0x00000002  0x6562      926      0/0/0
Metric: 0      IS-Extended P1.00
Metric: 0      IS-Extended PE1.00
P2.00-00      0x00000007  0xF048      1184      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P2
IP Address: 3.3.3.3
Router ID: 3.3.3.3
Metric: 10      IS-Extended P2.02
IPv4 Interface Address: 20.20.20.2
Neighbor IP Address: 20.20.20.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P2.04
IPv4 Interface Address: 70.70.70.2
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.04
IPv4 Interface Address: 80.80.80.2
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec

```



```

    Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P2.03
IPv4 Interface Address: 50.50.50.1
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 3.3.3.3/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P2.02-00            0x00000004  0x7055          967          0/0/0
Metric: 0           IS-Extended P2.00
Metric: 0           IS-Extended PE1.00
P2.03-00            0x00000004  0xB907          1184         0/0/0
Metric: 0           IS-Extended P2.00
Metric: 0           IS-Extended PE2.00
P2.04-00            0x00000004  0x764C          967          0/0/0
Metric: 0           IS-Extended P2.00
Metric: 0           IS-Extended P1.00
P3.00-00            0x00000006  0x72F1          1184         0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P3
IP Address: 4.4.4.4
Router ID: 4.4.4.4
Metric: 10          IS-Extended P3.02
IPv4 Interface Address: 30.30.30.2
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.04
IPv4 Interface Address: 80.80.80.1
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:

```

```

Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.03
IPv4 Interface Address: 60.60.60.1
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 4.4.4.4/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P3.02-00      0x00000004  0x7152      1043      0/0/0
Metric: 0      IS-Extended P3.00
Metric: 0      IS-Extended PE1.00
P3.03-00      0x00000004  0xBA04      1184      0/0/0
Metric: 0      IS-Extended P3.00
Metric: 0      IS-Extended PE2.00
P3.04-00      0x00000004  0x8B34      1043      0/0/0
Metric: 0      IS-Extended P3.00
Metric: 0      IS-Extended P2.00
PE2.00-00      * 0x00000005  0x859A      1189      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: PE2
IP Address: 5.5.5.5
Router ID: 5.5.5.5
Metric: 10      IS-Extended P2.03
IPv4 Interface Address: 50.50.50.2
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.03
IPv4 Interface Address: 60.60.60.2
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:

```

```

Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10 IS-Extended Pl.02
IPv4 Interface Address: 40.40.40.2
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10 IP-Extended 5.5.5.5/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set

PE2#ping 1.1.1.1
Press CTRL+C to exit
PING 1.1.1.1 (1.1.1.1) 100(128) bytes of data.
108 bytes from 1.1.1.1: icmp_seq=1 ttl=63 time=0.445 ms
108 bytes from 1.1.1.1: icmp_seq=2 ttl=63 time=0.460 ms
108 bytes from 1.1.1.1: icmp_seq=3 ttl=63 time=0.509 ms

--- 1.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2072ms
rtt min/avg/max/mdev = 0.445/0.471/0.509/0.027 ms

```

LDP FRR with IS-IS Configuration

Now that NSM, ISIS and LDP are all configured, FRR for IGP and LDP can be enabled using below configurations.

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 with instance 1 |
| (config-router)#fast-reroute per-prefix level-1 proto ipv4 all | Configure LFA-FRR to calculate the available backup path for all L1 ipv4 prefixes learnt |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#fast-reroute | Configure LDP LFA-FRR to calculate the available backup path |

| | |
|----------------------|-------------------------------|
| (config-router)#exit | Exit router mode. |
| (config)#commit | Commit all the configurations |

Validation

PE1

```
PE1#show clns neighbors
```

```
Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P1 | ce46/1 | b86a.97f5.ead2 | Up | 6 | L1 | IS-IS |
| P3 | ce46/3 | e8c5.7a8c.f2c8 | Up | 6 | L1 | IS-IS |
| P2 | ce46/4 | e49d.731b.ecc5 | Up | 7 | L1 | IS-IS |

```
PE1#show clns neighbors detail
```

```
Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P1 | ce46/1 | b86a.97f5.ead2 | Up | 6 | L1 | IS-IS |

```
L1 Adjacency ID: 1
```

```
L2 Adjacency ID: 2
```

```
Uptime: 00:26:34
```

```
Area Address(es): 49.0001
```

```
IP Address(es): 10.10.10.2
```

```
Level-1 Protocols Supported: IPv4
```

```
IPV4 Bidirectional Forwarding Detection is enabled
```

```
Adjacency advertisement: Advertise
```

| | | | | | | |
|----|--------|----------------|----|---|----|-------|
| P3 | ce46/3 | e8c5.7a8c.f2c8 | Up | 6 | L1 | IS-IS |
|----|--------|----------------|----|---|----|-------|

```
L1 Adjacency ID: 1
```

```
L2 Adjacency ID: 2
```

```
Uptime: 00:54:38
```

```
Area Address(es): 49.0001
```

```
IP Address(es): 30.30.30.2
```

```
...skipping 1 line
```

```
IPV4 Bidirectional Forwarding Detection is enabled
```

```
Adjacency advertisement: Advertise
```

| | | | | | | |
|----|--------|----------------|----|---|----|-------|
| P2 | ce46/4 | e49d.731b.ecc5 | Up | 7 | L1 | IS-IS |
|----|--------|----------------|----|---|----|-------|

```
L1 Adjacency ID: 1
```

```
L2 Adjacency ID: 2
```

```
Uptime: 00:55:53
```

```
Area Address(es): 49.0001
```

```
IP Address(es): 20.20.20.2
```

```
Level-1 Protocols Supported: IPv4
```

```
IPV4 Bidirectional Forwarding Detection is enabled
```

```
Adjacency advertisement: Advertise
```

```
PE1#s database
```

```
^
```

```
% Invalid input detected at '^' marker.
```

```
PE1#show isis database verbose
```

```
Tag 1: VRF : default
```

```
IS-IS Level-1 Link State Database:
```

```

LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00            * 0x00000009  0xA301        508           0/0/0
  Area Address: 49.0001
  NLPID:         0xCC
  Hostname:      PE1
  IP Address:    1.1.1.1
  Router ID:     1.1.1.1
  Metric: 15     IS-Extended P2.02
    IPv4 Interface Address: 20.20.20.1
    Neighbor IP Address: 20.20.20.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 15
  Metric: 60     IS-Extended P3.02
    IPv4 Interface Address: 30.30.30.1
    Neighbor IP Address: 30.30.30.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 60
  Metric: 20     IS-Extended P1.04
    IPv4 Interface Address: 10.10.10.1
    Neighbor IP Address: 10.10.10.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 20
  Metric: 10     IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric: 60     IP-Extended 30.30.30.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric: 15     IP-Extended 20.20.20.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric: 20     IP-Extended 10.10.10.0/24
    Prefix Attribute Flags[0]: ELC Set
P1.00-00            0x00000007  0xFB9F        511           0/0/0
  Area Address: 49.0001
  NLPID:         0xCC
  Hostname:      P1
  IP Address:    2.2.2.2
  Router ID:     2.2.2.2
  Metric: 10     IS-Extended P2.04
    IPv4 Interface Address: 70.70.70.1

```

```

Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec

...skipping 1 line
Metric: 10          IS-Extended P1.02
IPv4 Interface Address: 40.40.40.1
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P1.04
IPv4 Interface Address: 10.10.10.2
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 2.2.2.2/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
P1.02-00            0x00000004  0xBF04          765          0/0/0
Metric: 0           IS-Extended P1.00
Metric: 0           IS-Extended PE2.00
P1.04-00            0x00000002  0x6562          507          0/0/0
Metric: 0           IS-Extended P1.00
Metric: 0           IS-Extended PE1.00
P2.00-00            0x00000007  0xF048          765          0/0/0
Area Address: 49.0001
NLPID:              0xCC

...skipping 23 lines
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec

```

```

    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.04
IPv4 Interface Address: 80.80.80.2
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P2.03
IPv4 Interface Address: 50.50.50.1
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 3.3.3.3/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P2.02-00      0x00000004  0x7055      548      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended PE1.00
P2.03-00      0x00000004  0xB907      765      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended PE2.00
P2.04-00      0x00000004  0x764C      548      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended P1.00
P3.00-00      0x00000006  0x72F1      765      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P3
IP Address: 4.4.4.4
Router ID: 4.4.4.4
Metric: 10      IS-Extended P3.02
IPv4 Interface Address: 30.30.30.2
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec

...skipping 1 line
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec

```

```

    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.04
IPv4 Interface Address: 80.80.80.1
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.03
IPv4 Interface Address: 60.60.60.1
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 4.4.4.4/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
Most commands optionally preceded by integer argument k.  Defaults in brackets.
Star (*) indicates argument becomes new default.
-----
<space>          Display next k lines of text [current screen size]
z                Display next k lines of text [current screen size]*
<return>         Display next k lines of text [1]*
d or ctrl-D      Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s                Skip forward k lines of text [1]
f                Skip forward k screenfuls of text [1]
b or ctrl-B      Skip backwards k screenfuls of text [1]
'                Go to place where previous search started
=                Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n                Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v                Start up '/usr/bin/vi' at current line
ctrl-L           Redraw screen
:n               Go to kth next file [1]
:p               Go to kth previous file [1]
:f               Display current file name and line number
.               Repeat previous command

```



```

-----
P3.02-00          0x00000004  0x7152      624          0/0/0
Metric:   0          IS-Extended P3.00
Metric:   0          IS-Extended PE1.00
P3.03-00          0x00000004  0xBA04      765          0/0/0
Metric:   0          IS-Extended P3.00
Metric:   0          IS-Extended PE2.00
P3.04-00          0x00000004  0x8B34      624          0/0/0
Metric:   0          IS-Extended P3.00
Metric:   0          IS-Extended P2.00
PE2.00-00          0x00000005  0x859A      769          0/0/0
Area Address: 49.0001

...skipping 1 line
Hostname:      PE2
IP Address:    5.5.5.5
Router ID:     5.5.5.5
Metric:   10          IS-Extended P2.03
IPv4 Interface Address: 50.50.50.2
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric:   10          IS-Extended P3.03
IPv4 Interface Address: 60.60.60.2
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric:   10          IS-Extended P1.02
Most commands optionally preceded by integer argument k.  Defaults in brackets.
Star (*) indicates argument becomes new default.
-----
<space>          Display next k lines of text [current screen size]
z                Display next k lines of text [current screen size]*
<return>         Display next k lines of text [1]*
d or ctrl-D      Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s                Skip forward k lines of text [1]
f                Skip forward k screenfuls of text [1]
b or ctrl-B      Skip backwards k screenfuls of text [1]
'                Go to place where previous search started
=                Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n                Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v                Start up '/usr/bin/vi' at current line
ctrl-L           Redraw screen
:n               Go to kth next file [1]
:p               Go to kth previous file [1]

```

```

:f          Display current file name and line number
.          Repeat previous command
-----
IPv4 Interface Address: 40.40.40.2
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 5.5.5.5/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set

PE1#forwarding-table
^
% Invalid input detected at '^' marker.

PE1#show mpls forwarding-table 5.5.5.5/32
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code  FEC          Nexthop    FTN-ID  Nhlfe-ID  Tunnel-ID  Pri  Out-Label  Out-
Intf   ELC           N/A      Algo-Num UpTime
L>    5.5.5.5/32    2        25      -          -      -      -          -
      -           N/A      01:39:16
      4           -
      Yes  24324      ce46/4    No       20.20.20.2  -      -
      24           -
      No   24322      ce46/1    No       10.10.10.2  -      -

PE1#show mpls ftn-table
Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:39:16, UpTime: 01:39:16, LastUpdate: 00:03:35
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 15 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 15, owner: N/A, Stale: NO, refcount: 6, out intf: ce46/1, out label: 3
Nexthop addr: 10.10.10.2      cross connect ix: 13, op code: Push

Backup Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: LDP, Stale: NO, refcount: 5, out intf: ce46/4, out label: 24323
Nexthop addr: 20.20.20.2      cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:39:16, UpTime: 01:39:16, LastUpdate: 00:03:35
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0

```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
  Nexthop addr: 20.20.20.2      cross connect ix: 6, op code: Push

  Backup Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 5 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 5, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label: 24320
  Nexthop addr: 10.10.10.2      cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:39:16, UpTime: 01:39:16, LastUpdate: 00:03:35
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 62 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 62, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24320
  Nexthop addr: 20.20.20.2      cross connect ix: 11, op code: Push

  Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 21 bypass ftn-ix: 0
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 21, owner: N/A, Stale: NO, refcount: 10, out intf: ce46/3, out label: 3
  Nexthop addr: 30.30.30.2      cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:39:16, UpTime: 01:39:16, LastUpdate: 00:03:35
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 4, owner: LDP, Stale: NO, refcount: 4, out intf: ce46/4, out label: 24324
  Nexthop addr: 20.20.20.2      cross connect ix: 3, op code: Push

  Backup Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 24 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 24, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24322
  Nexthop addr: 10.10.10.2      cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 40.40.40.0/24, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:39:16, UpTime: 01:39:16, LastUpdate: 00:03:35
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 15 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 15, owner: N/A, Stale: NO, refcount: 6, out intf: ce46/1, out label: 3
  Nexthop addr: 10.10.10.2      cross connect ix: 13, op code: Push

  Backup Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 26 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 26, owner: LDP, Stale: NO, refcount: 4, out intf: ce46/3, out label:
24326
  Nexthop addr: 30.30.30.2      cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 50.50.50.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:39:16, UpTime: 01:39:16, LastUpdate: 00:03:35
```

```
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1

...skipping 1 line
Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 6, op code: Push

Backup Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 33 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 33, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24325
Nexthop addr: 10.10.10.2          cross connect ix: 10, op code: Push

Primary FTN entry with FEC: 60.60.60.0/24, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:39:16, UpTime: 01:39:16, LastUpdate: 00:03:35
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 72 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 72, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24321
Nexthop addr: 20.20.20.2          cross connect ix: 12, op code: Push

Backup Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 37 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 37, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/1, out label:
24326
Nexthop addr: 10.10.10.2          cross connect ix: 12, op code: Push

...skipping 1 line
Primary FTN entry with FEC: 70.70.70.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:39:16, UpTime: 01:39:16, LastUpdate: 00:03:35
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 6, op code: Push

Backup Cross connect ix: 14, in intf: - in label: 0 out-segment ix: 15 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 15, owner: N/A, Stale: NO, refcount: 6, out intf: ce46/1, out label: 3
Nexthop addr: 10.10.10.2          cross connect ix: 13, op code: Push

Primary FTN entry with FEC: 80.80.80.0/24, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:39:16, UpTime: 01:39:16, LastUpdate: 00:03:35
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
Nexthop addr: 20.20.20.2          cross connect ix: 6, op code: Push

Backup Cross connect ix: 16, in intf: - in label: 0 out-segment ix: 21 bypass ftn-ix: 0
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 21, owner: N/A, Stale: NO, refcount: 10, out intf: ce46/3, out label: 3
Nexthop addr: 30.30.30.2          cross connect ix: 6, op code: Push
```

```

PE1#clns neighbors detail
^
% Invalid input detected at '^' marker.

PE1#show isis database
Tag 1: VRF : default
IS-IS Level-1 Link State Database:
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00             * 0x00000009  0xA301        505           0/0/0
P1.00-00              0x00000007  0xFB9F        508           0/0/0
P1.02-00              0x00000004  0xBF04        762           0/0/0
P1.04-00              0x00000002  0x6562        504           0/0/0
P2.00-00              0x00000007  0xF048        762           0/0/0
P2.02-00              0x00000004  0x7055        545           0/0/0
P2.03-00              0x00000004  0xB907        762           0/0/0
P2.04-00              0x00000004  0x764C        545           0/0/0
P3.00-00              0x00000006  0x72F1        762           0/0/0
P3.02-00              0x00000004  0x7152        621           0/0/0
P3.03-00              0x00000004  0xBA04        762           0/0/0
P3.04-00              0x00000004  0x8B34        621           0/0/0
PE2.00-00             0x00000005  0x859A        766           0/0/0

PE1#show isis database verbose
Tag 1: VRF : default
IS-IS Level-1 Link State Database:
LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00             * 0x00000009  0xA301        505           0/0/0
Area Address: 49.0001
NLPID:               0xCC
Hostname:            PE1
IP Address:          1.1.1.1
Router ID:           1.1.1.1
Metric: 15           IS-Extended P2.02
IPv4 Interface Address: 20.20.20.1
Neighbor IP Address: 20.20.20.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 15
Metric: 60           IS-Extended P3.02
IPv4 Interface Address: 30.30.30.1
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 60
Metric: 20           IS-Extended P1.04
IPv4 Interface Address: 10.10.10.1
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec

```

```

Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 20
Metric: 10      IP-Extended 1.1.1.1/32
  Prefix Attribute Flags[0]: ELC Set
Metric: 60      IP-Extended 30.30.30.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric: 15      IP-Extended 20.20.20.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric: 20      IP-Extended 10.10.10.0/24
  Prefix Attribute Flags[0]: ELC Set
Pl.00-00      0x00000007  0xFB9F      508      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P1
IP Address: 2.2.2.2
Router ID: 2.2.2.2
Metric: 10      IS-Extended P2.04
  IPv4 Interface Address: 70.70.70.1
  Neighbor IP Address: 70.70.70.2
  Maximum Link Bandwidth : 10000000.00 kbits/sec
  Reservable Bandwidth : 10000000.00 kbits/sec
  Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
  TE-Default Metric: 10
Metric: 10      IS-Extended P1.02
  IPv4 Interface Address: 40.40.40.1
  Neighbor IP Address: 40.40.40.1
  Maximum Link Bandwidth : 10000000.00 kbits/sec
  Reservable Bandwidth : 10000000.00 kbits/sec
  Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
  TE-Default Metric: 10
Metric: 10      IS-Extended P1.04
  IPv4 Interface Address: 10.10.10.2
  Neighbor IP Address: 10.10.10.2
  Maximum Link Bandwidth : 10000000.00 kbits/sec
  Reservable Bandwidth : 10000000.00 kbits/sec
  Unreserved Bandwidth:
    Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec

```

```

TE-Default Metric: 10
Metric: 10 IP-Extended 2.2.2.2/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
P1.02-00 0x00000004 0xBF04 762 0/0/0
Metric: 0 IS-Extended P1.00
Metric: 0 IS-Extended PE2.00
P1.04-00 0x00000002 0x6562 504 0/0/0
Metric: 0 IS-Extended P1.00
Metric: 0 IS-Extended PE1.00
P2.00-00 0x00000007 0xF048 762 0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P2
IP Address: 3.3.3.3
Router ID: 3.3.3.3
Metric: 10 IS-Extended P2.02
IPv4 Interface Address: 20.20.20.2
Neighbor IP Address: 20.20.20.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10 IS-Extended P2.04
IPv4 Interface Address: 70.70.70.2
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10 IS-Extended P3.04
IPv4 Interface Address: 80.80.80.2
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:
Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10 IS-Extended P2.03
IPv4 Interface Address: 50.50.50.1

```

```

Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 3.3.3.3/32
  Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 20.20.20.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 50.50.50.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 70.70.70.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 80.80.80.0/24
  Prefix Attribute Flags[0]: ELC Set
P2.02-00      0x00000004  0x7055      545      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended PE1.00
P2.03-00      0x00000004  0xB907      762      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended PE2.00
P2.04-00      0x00000004  0x764C      545      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended P1.00
P3.00-00      0x00000006  0x72F1      762      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P3
IP Address: 4.4.4.4
Router ID: 4.4.4.4
Metric: 10      IS-Extended P3.02
IPv4 Interface Address: 30.30.30.2
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P3.04
IPv4 Interface Address: 80.80.80.1
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10

```



```

Metric: 10          IS-Extended P3.03
IPv4 Interface Address: 60.60.60.1
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 4.4.4.4/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P3.02-00            0x00000004 0x7152          621          0/0/0
Metric: 0           IS-Extended P3.00
Metric: 0           IS-Extended PE1.00
P3.03-00            0x00000004 0xBA04          762          0/0/0
Metric: 0           IS-Extended P3.00
Metric: 0           IS-Extended PE2.00
P3.04-00            0x00000004 0x8B34          621          0/0/0
Metric: 0           IS-Extended P3.00
Metric: 0           IS-Extended P2.00
PE2.00-00           0x00000005 0x859A          766          0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: PE2
IP Address: 5.5.5.5
Router ID: 5.5.5.5
Metric: 10          IS-Extended P2.03
IPv4 Interface Address: 50.50.50.2
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.03
IPv4 Interface Address: 60.60.60.2
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10

```

```

Metric: 10          IS-Extended P1.02
IPv4 Interface Address: 40.40.40.2
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 5.5.5.5/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set

PE1#ping 5.5.5.5
Press CTRL+C to exit
PING 5.5.5.5 (5.5.5.5) 100(128) bytes of data.
108 bytes from 5.5.5.5: icmp_seq=1 ttl=64 time=0.602 ms
108 bytes from 5.5.5.5: icmp_seq=2 ttl=64 time=0.508 ms
108 bytes from 5.5.5.5: icmp_seq=3 ttl=64 time=0.624 ms
108 bytes from 5.5.5.5: icmp_seq=4 ttl=64 time=0.384 ms
108 bytes from 5.5.5.5: icmp_seq=5 ttl=64 time=0.484 ms

--- 5.5.5.5 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4086ms
rtt min/avg/max/mdev = 0.384/0.520/0.624/0.086 ms

```

Backup Path based on Route-Map Prefixes

| | |
|---|--|
| (config)#ip access-list A | Create an access-list |
| (config-ip-acl)#10 permit any host 40.40.40.0 any | Configuring rule to permit only one prefix |
| (config)#route-map R permit 10 | Create a route-map |
| (config-route-map)#match ip address A | Apply above created access-list in route-map |
| (config-route-map)#exit | Exit route-map mode |
| (config)#commit | Commit all the configurations |

Apply the above created route-map with fast-reroute:

| | |
|--|---|
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 with instance 1 |
| (config-router)#no fast-reroute per-prefix level-1 proto ipv4 | Un-configure LFA-FRR to calculate available path for all prefixes |
| (config-router)#fast-reroute per-prefix level-1 proto ipv4 route-map R | Configure LFA-FRR to calculate the available backup path for routes allowed through route-map |

| | |
|----------------------|-------------------------------|
| (config-router)#exit | Exit router mode |
| (config)#commit | Commit all the configurations |

Validation

PE1

```
PE1#show clns neighbors
```

```
Total number of L1 adjacencies: 3
```

```
Total number of L2 adjacencies: 0
```

```
Total number of adjacencies: 3
```

```
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P1 | ce46/1 | b86a.97f5.ead2 | Up | 5 | L1 | IS-IS |
| P3 | ce46/3 | e8c5.7a8c.f2c8 | Up | 6 | L1 | IS-IS |
| P2 | ce46/4 | e49d.731b.ecc5 | Up | 7 | L1 | IS-IS |

```
PE1#show clns neighbors detail
```

```
Total number of L1 adjacencies: 3
```

```
Total number of L2 adjacencies: 0
```

```
Total number of adjacencies: 3
```

```
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P1 | ce46/1 | b86a.97f5.ead2 | Up | 5 | L1 | IS-IS |

```
L1 Adjacency ID: 1
```

```
L2 Adjacency ID: 2
```

```
Uptime: 00:39:54
```

```
Area Address(es): 49.0001
```

```
IP Address(es): 10.10.10.2
```

```
Level-1 Protocols Supported: IPv4
```

```
IPV4 Bidirectional Forwarding Detection is enabled
```

```
Adjacency advertisement: Advertise
```

| | | | | | | |
|----|--------|----------------|----|---|----|-------|
| P3 | ce46/3 | e8c5.7a8c.f2c8 | Up | 6 | L1 | IS-IS |
|----|--------|----------------|----|---|----|-------|

```
L1 Adjacency ID: 1
```

```
L2 Adjacency ID: 2
```

```
Uptime: 01:07:58
```

```
Area Address(es): 49.0001
```

```
IP Address(es): 30.30.30.2
```

```
...skipping 1 line
```

```
IPV4 Bidirectional Forwarding Detection is enabled
```

```
Adjacency advertisement: Advertise
```

| | | | | | | |
|----|--------|----------------|----|---|----|-------|
| P2 | ce46/4 | e49d.731b.ecc5 | Up | 7 | L1 | IS-IS |
|----|--------|----------------|----|---|----|-------|

```
L1 Adjacency ID: 1
```

```
L2 Adjacency ID: 2
```

```
Uptime: 01:09:13
```

```
Area Address(es): 49.0001
```

```
IP Address(es): 20.20.20.2
```

```
Level-1 Protocols Supported: IPv4
```

```
IPV4 Bidirectional Forwarding Detection is enabled
```

```
Adjacency advertisement: Advertise
```

```
PE1#s database
```

```
^
```

```
% Invalid input detected at '^' marker.
```

```
PE1#show isis database verbose
```

```
Tag 1: VRF : default
```

```
IS-IS Level-1 Link State Database:
```

```

LSPID                LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00            * 0x0000000A  0xA102        607           0/0/0
  Area Address: 49.0001
  NLPID:          0xCC
  Hostname:       PE1
  IP Address:     1.1.1.1
  Router ID:      1.1.1.1
  Metric: 15      IS-Extended P2.02
    IPv4 Interface Address: 20.20.20.1
    Neighbor IP Address: 20.20.20.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 15
  Metric: 60      IS-Extended P3.02
    IPv4 Interface Address: 30.30.30.1
    Neighbor IP Address: 30.30.30.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
...skipping 1 line
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 60
  Metric: 20      IS-Extended P1.04
    IPv4 Interface Address: 10.10.10.1
    Neighbor IP Address: 10.10.10.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 20
  Metric: 10      IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric: 60      IP-Extended 30.30.30.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric: 15      IP-Extended 20.20.20.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric: 20      IP-Extended 10.10.10.0/24
    Prefix Attribute Flags[0]: ELC Set
Pl.00-00             0x00000008  0xF9A0        610           0/0/0
  Area Address: 49.0001
  NLPID:          0xCC
  Hostname:       P1
  IP Address:     2.2.2.2
  Router ID:      2.2.2.2
  Metric: 10      IS-Extended P2.04

```

```
IPv4 Interface Address: 70.70.70.1
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
```

Most commands optionally preceded by integer argument k. Defaults in brackets.
 Star (*) indicates argument becomes new default.

```
-----
<space>          Display next k lines of text [current screen size]
z                Display next k lines of text [current screen size]*
<return>         Display next k lines of text [1]*
d or ctrl-D      Scroll k lines [current scroll size, initially 11]*
q or Q or <interrupt> Exit from more
s                Skip forward k lines of text [1]
f                Skip forward k screenfuls of text [1]
b or ctrl-B      Skip backwards k screenfuls of text [1]
'                Go to place where previous search started
=                Display current line number
/<regular expression> Search for kth occurrence of regular expression [1]
n                Search for kth occurrence of last r.e [1]
!<cmd> or :!<cmd> Execute <cmd> in a subshell
v                Start up '/usr/bin/vi' at current line
ctrl-L           Redraw screen
:n               Go to kth next file [1]
:p               Go to kth previous file [1]
:f               Display current file name and line number
.               Repeat previous command
-----
```

```
Metric: 10      IS-Extended Pl.02
IPv4 Interface Address: 40.40.40.1
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended Pl.04
IPv4 Interface Address: 10.10.10.2
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
...skipping 1 line
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
```

```

Metric: 10      IP-Extended 2.2.2.2/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
P1.02-00      0x00000005  0xBD05      864      0/0/0
Metric: 0      IS-Extended P1.00
Metric: 0      IS-Extended PE2.00
P1.04-00      0x00000003  0x6363      606      0/0/0
Metric: 0      IS-Extended P1.00
Metric: 0      IS-Extended PE1.00
P2.00-00      0x00000008  0xEE49      865      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P2
IP Address: 3.3.3.3
Router ID: 3.3.3.3
Metric: 10      IS-Extended P2.02
IPv4 Interface Address: 20.20.20.2
Neighbor IP Address: 20.20.20.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P2.04
IPv4 Interface Address: 70.70.70.2
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec

...skipping 23 lines
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P2.03
IPv4 Interface Address: 50.50.50.1
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IP-Extended 3.3.3.3/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 80.80.80.0/24

```

```

Prefix Attribute Flags[0]: ELC Set
P2.02-00      0x00000005 0x6E56      647      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended PE1.00
P2.03-00      0x00000005 0xB708      865      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended PE2.00
P2.04-00      0x00000005 0x744D      647      0/0/0
Metric: 0      IS-Extended P2.00
Metric: 0      IS-Extended P1.00
P3.00-00      0x00000007 0x70F2      865      0/0/0
Area Address: 49.0001
NLPID:        0xCC
Hostname:      P3
IP Address:    4.4.4.4
Router ID:     4.4.4.4
Metric: 10     IS-Extended P3.02
IPv4 Interface Address: 30.30.30.2
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10     IS-Extended P3.04
IPv4 Interface Address: 80.80.80.1
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
...skipping 1 line
  Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10     IS-Extended P3.03
IPv4 Interface Address: 60.60.60.1
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10     IP-Extended 4.4.4.4/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10     IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10     IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set

```

```

Metric: 10          IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P3.02-00          0x00000005 0x6F53          723          0/0/0
Metric: 0          IS-Extended P3.00
Metric: 0          IS-Extended PE1.00
P3.03-00          0x00000005 0xB805          865          0/0/0
Metric: 0          IS-Extended P3.00
Metric: 0          IS-Extended PE2.00
P3.04-00          0x00000005 0x8935          723          0/0/0
Metric: 0          IS-Extended P3.00
Metric: 0          IS-Extended P2.00
PE2.00-00          0x00000006 0x839B          868          0/0/0
Area Address: 49.0001
NLPID:            0xCC
Hostname:         PE2
IP Address:       5.5.5.5
Router ID:        5.5.5.5
Metric: 10          IS-Extended P2.03
IPv4 Interface Address: 50.50.50.2
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.03
IPv4 Interface Address: 60.60.60.2
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P1.02
IPv4 Interface Address: 40.40.40.2
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 5.5.5.5/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set

```



```

Metric: 10 IP-Extended 40.40.40.0/24

...skipping 1 line

PE1#show ldp downstream
^
% Invalid input detected at '^' marker.

PE1#show ldp lsp
DOWNSTREAM LSP :

```

| FEC | State | NextHop Label | Req.ID | Attr | Code | | | |
|---------------|-------------|---------------|-------------|-------|------|------|----|----|
| 1.1.1.1/32 | connected | | Established | 24328 | 0 | None | | No |
| 1.1.1.1/32 | connected | | Established | 24327 | 0 | None | | No |
| 1.1.1.1/32 | connected | | Established | 24326 | 0 | None | | No |
| 1.1.1.1/32 | connected | | Established | none | 0 | None | | No |
| 1.1.1.1/32 | connected | | Established | 24327 | 0 | None | | No |
| 2.2.2.2/32 | 10.10.10.2 | Established | impl- | | | | | |
| null 0 | None | Yes | P | | | | | |
| 2.2.2.2/32 | non- | | | | | | | |
| nh | Established | 24324 | 0 | None | No | - | | |
| 2.2.2.2/32 | non- | | | | | | | |
| nh | Established | 24323 | 0 | None | No | - | | |
| 2.2.2.2/32 | non- | | | | | | | |
| nh | Established | 24320 | 0 | None | No | - | | |
| 3.3.3.3/32 | non- | | | | | | | |
| nh | Established | 24320 | 0 | None | No | - | | |
| 3.3.3.3/32 | non- | | | | | | | |
| nh | Established | 24320 | 0 | None | No | - | | |
| 3.3.3.3/32 | 20.20.20.2 | Established | impl- | | | | | |
| null 0 | None | Yes | P | | | | | |
| 3.3.3.3/32 | non- | | | | | | | |
| nh | Established | 24321 | 0 | None | No | - | | |
| 4.4.4.4/32 | non- | | | | | | | |
| nh | Established | 24321 | 0 | None | No | - | | |
| 4.4.4.4/32 | 20.20.20.2 | Established | 24320 | 0 | None | | Ye | |
| s P | | | | | | | | |
| 4.4.4.4/32 | non-nh | Established | impl- | | | | | |
| null 0 | None | No | - | | | | | |
| 4.4.4.4/32 | non- | | | | | | | |
| nh | Established | 24322 | 0 | None | No | - | | |
| 5.5.5.5/32 | non- | | | | | | | |
| nh | Established | 24322 | 0 | None | No | - | | |
| 5.5.5.5/32 | non- | | | | | | | |
| nh | Established | 24325 | 0 | None | No | - | | |
| 5.5.5.5/32 | 20.20.20.2 | Established | 24324 | 0 | None | | Ye | |
| s P | | | | | | | | |
| 5.5.5.5/32 | non- | | | | | | | |
| nh | Established | 24328 | 0 | None | No | - | | |
| 10.10.10.0/24 | connected | Established | impl- | | | | | |
| null 0 | None | No | - | | | | | |
| 10.10.10.0/24 | connected | Established | 24328 | 0 | None | | No | |
| - | | | | | | | | |
| 10.10.10.0/24 | connected | Established | 24327 | 0 | None | | No | |
| - | | | | | | | | |
| 10.10.10.0/24 | connected | Established | none | 0 | None | | No | |
| - | | | | | | | | |
| 20.20.20.0/24 | connected | Established | 24323 | 0 | None | | No | |
| - | | | | | | | | |
| 20.20.20.0/24 | connected | Established | impl- | | | | | |
| null 0 | None | No | - | | | | | |
| 20.20.20.0/24 | connected | Established | 24323 | 0 | None | | No | |
| - | | | | | | | | |
| 20.20.20.0/24 | connected | Established | 24323 | 0 | None | | No | |
| - | | | | | | | | |
| 20.20.20.0/24 | connected | Established | none | 0 | None | | No | |

| | | | | | | |
|----------------------|-------------|-------------|-------|------|---------|--------|
| 30.30.30.0/24 | connected | Established | 24324 | 0 | None | No |
| 30.30.30.0/24 | connected | Established | impl- | | | |
| null 0 | None | No | - | | | |
| 30.30.30.0/24 | connected | Established | 24322 | 0 | None | No |
| 30.30.30.0/24 | connected | Established | 24324 | 0 | None | No |
| 30.30.30.0/24 | connected | Established | none | 0 | None | No |
| 40.40.40.0/24 | 10.10.10.2 | Established | impl- | | | |
| null 0 | None | Yes | P | | | |
| 40.40.40.0/24 | 30.30.30.2 | Established | 24326 | 0 | None | Ye |
| s 40.40.40.0/24 | non- | | | | | |
| nh | Established | 24325 | 0 | None | No | - |
| 40.40.40.0/24 | non- | | | | | |
| nh | Established | 24329 | 0 | None | No | - |
| 50.50.50.0/24 | non- | | | | | |
| nh | Established | 24325 | 0 | None | No | - |
| 50.50.50.0/24 | non- | | | | | |
| nh | Established | 24321 | 0 | None | No | - |
| 50.50.50.0/24 | 20.20.20.2 | Established | impl- | | | |
| null 0 | None | Yes | P | | | |
| 50.50.50.0/24 | non- | | | | | |
| nh | Established | 24330 | 0 | None | No | - |
| 60.60.60.0/24 | non- | | | | | |
| nh | Established | 24326 | 0 | None | No | - |
| ...skipping 23 lines | | | | | | |
| 3.3.3.3/32 | Established | 24322 | 0 | None | Yes | /yes 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | skipped | /no 4 |
| 3.3.3.3/32 | Established | 24322 | 0 | None | Yes | /yes 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | Yes | /yes 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | Yes | /yes 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | skipped | /no 4 |
| 4.4.4.4/32 | Established | 24323 | 0 | None | Yes | /yes 4 |
| 5.5.5.5/32 | Established | 24324 | 0 | None | Yes | /yes 4 |
| 5.5.5.5/32 | Established | 24324 | 0 | None | Yes | /yes 4 |
| 5.5.5.5/32 | Established | 24324 | 0 | None | skipped | /no 4 |
| 5.5.5.5/32 | Established | 24324 | 0 | None | Yes | /yes 4 |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 10.10.10.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 20.20.20.0/24 | Established | 24325 | 0 | None | Yes | /yes 1 |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 20.20.20.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 30.30.30.0/24 | Established | 24326 | 0 | None | Yes | /yes 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 30.30.30.0/24 | Established | impl-null | 0 | None | No | /yes 1 |
| 40.40.40.0/24 | Established | 24327 | 0 | None | Yes | /yes 4 |
| 40.40.40.0/24 | Established | 24327 | 0 | None | skipped | /no 4 |
| 40.40.40.0/24 | Established | 24327 | 0 | None | Yes | /yes 4 |
| 40.40.40.0/24 | Established | 24327 | 0 | None | Yes | /yes 4 |
| 40.40.40.0/24 | Established | 24327 | 0 | None | Yes | /yes 4 |
| 40.40.40.0/24 | Established | 24327 | 0 | None | skipped | /no 4 |
| 40.40.40.0/24 | Established | 24327 | 0 | None | Yes | /yes 4 |
| 40.40.40.0/24 | Established | 24327 | 0 | None | Yes | /yes 4 |
| 50.50.50.0/24 | Established | 24328 | 0 | None | skipped | /no 4 |
| 50.50.50.0/24 | Established | 24328 | 0 | None | Yes | /yes 4 |
| 50.50.50.0/24 | Established | 24328 | 0 | None | Yes | /yes 4 |
| 50.50.50.0/24 | Established | 24328 | 0 | None | Yes | /yes 4 |
| 50.50.50.0/24 | Established | 24328 | 0 | None | Yes | /yes 4 |
| 60.60.60.0/24 | Established | 24329 | 0 | None | Yes | /yes 4 |
| 60.60.60.0/24 | Established | 24329 | 0 | None | Yes | /yes 4 |
| 60.60.60.0/24 | Established | 24329 | 0 | None | skipped | /no 4 |
| 60.60.60.0/24 | Established | 24329 | 0 | None | Yes | /yes 4 |

```

70.70.70.0/24      Established      24330      0      None      Yes      /yes      4
70.70.70.0/24      Established      24330      0      None      Yes      /yes      4
70.70.70.0/24      Established      24330      0      None      skipped  /no      4
70.70.70.0/24      Established      24330      0      None      Yes      /yes      4
80.80.80.0/24      Established      24331      0      None      Yes      /yes      4
80.80.80.0/24      Established      24331      0      None      Yes      /yes      4
80.80.80.0/24      Established      24331      0      None      skipped  /no      4
80.80.80.0/24      Established      24331      0      None      Yes      /yes      4
PE1#mpls forwarding-table 5.5.5.5/32
^
% Invalid input detected at '^' marker.

PE1#show mpls ftn-table
Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:52:37, UpTime: 01:52:37, LastUpdate: 00:03:59
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 15 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 15, owner: N/A, Stale: NO, refcount: 4, out intf: ce46/1, out label: 3
  Nexthop addr: 10.10.10.2      cross connect ix: 13, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:52:37, UpTime: 01:52:37, LastUpdate: 00:03:59
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
  Nexthop addr: 20.20.20.2      cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:52:37, UpTime: 01:52:37, LastUpdate: 00:03:59
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 62 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 62, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24320
  Nexthop addr: 20.20.20.2      cross connect ix: 11, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:52:37, UpTime: 01:52:37, LastUpdate: 00:03:59
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 4, owner: LDP, Stale: NO, refcount: 4, out intf: ce46/4, out label: 24324
  Nexthop addr: 20.20.20.2      cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 40.40.40.0/24, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:52:37, UpTime: 01:52:37, LastUpdate: 00:03:51
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 15 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

```

```

    Out-segment with ix: 15, owner: N/A, Stale: NO, refcount: 4, out intf: ce46/1, out label: 3
    Nexthop addr: 10.10.10.2          cross connect ix: 13, op code: Push

    Backup Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 26 bypass ftn-ix: 0
    Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 26, owner: LDP, Stale: NO, refcount: 4, out intf: ce46/3, out label:
24326
    Nexthop addr: 30.30.30.2          cross connect ix: 4, op code: Push

    Primary FTN entry with FEC: 50.50.50.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
    CreateTime: 01:52:37, UpTime: 01:52:37, LastUpdate: 00:03:59
    Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
    Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
    Nexthop addr: 20.20.20.2          cross connect ix: 6, op code: Push

    Primary FTN entry with FEC: 60.60.60.0/24, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
    CreateTime: 01:52:37, UpTime: 01:52:37, LastUpdate: 00:03:59
    Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
    Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 12, in intf: - in label: 0 out-segment ix: 72 refcount: 1
    Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 72, owner: LDP, Stale: NO, refcount: 2, out intf: ce46/4, out label:
24321
    Nexthop addr: 20.20.20.2          cross connect ix: 12, op code: Push

    Primary FTN entry with FEC: 70.70.70.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
    CreateTime: 01:52:37, UpTime: 01:52:37, LastUpdate: 00:03:59
    Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
    Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1

...skipping 1 line
    Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
    Nexthop addr: 20.20.20.2          cross connect ix: 6, op code: Push

    Primary FTN entry with FEC: 80.80.80.0/24, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
    CreateTime: 01:52:37, UpTime: 01:52:37, LastUpdate: 00:03:59
    Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
    Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 16 refcount: 1
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 13, out intf: ce46/4, out label: 3
    Nexthop addr: 20.20.20.2          cross connect ix: 6, op code: Push

PE1#show clns neighbors
^
% Invalid input detected at '^' marker.

PE1#show clns neighbors detail

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3

```

```

Tag 1: VRF : default
System Id      Interface      SNPA      State  Holdtime  Type  Protocol
P1             ce46/1      b86a.97f5.ead2  Up    5          L1    IS-IS
  L1 Adjacency ID: 1
  L2 Adjacency ID: 2
  Uptime: 00:39:58
  Area Address(es): 49.0001
  IP Address(es): 10.10.10.2
  Level-1 Protocols Supported: IPv4
  IPV4 Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

P3             ce46/3      e8c5.7a8c.f2c8  Up    6          L1    IS-IS
  L1 Adjacency ID: 1
  L2 Adjacency ID: 2
  Uptime: 01:08:02
  Area Address(es): 49.0001
  IP Address(es): 30.30.30.2

...skipping 1 line
  IPV4 Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

P2             ce46/4      e49d.731b.ecc5  Up    7          L1    IS-IS
  L1 Adjacency ID: 1
  L2 Adjacency ID: 2
  Uptime: 01:09:17
  Area Address(es): 49.0001
  IP Address(es): 20.20.20.2
  Level-1 Protocols Supported: IPv4
  IPV4 Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise

PE1#s database
      ^
% Invalid input detected at '^' marker.

PE1#show isis database verbose
Tag 1: VRF : default
IS-IS Level-1 Link State Database:
LSPID      LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
PE1.00-00   * 0x0000000A  0xA102        604           0/0/0
  Area Address: 49.0001
  NLPID:      0xCC
  Hostname:    PE1
  IP Address:  1.1.1.1
  Router ID:   1.1.1.1
  Metric:      15      IS-Extended P2.02
    IPv4 Interface Address: 20.20.20.1
    Neighbor IP Address: 20.20.20.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
      Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
      Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 15
  Metric:      60      IS-Extended P3.02
    IPv4 Interface Address: 30.30.30.1
    Neighbor IP Address: 30.30.30.2

```

```

Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 60
Metric: 20      IS-Extended P1.04
IPv4 Interface Address: 10.10.10.1
Neighbor IP Address: 10.10.10.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 20
Metric: 10      IP-Extended 1.1.1.1/32
Prefix Attribute Flags[0]: ELC Set
Metric: 60      IP-Extended 30.30.30.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 15      IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 20      IP-Extended 10.10.10.0/24
Prefix Attribute Flags[0]: ELC Set
P1.00-00      0x00000008  0xF9A0      607      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P1
IP Address: 2.2.2.2
Router ID: 2.2.2.2
Metric: 10      IS-Extended P2.04
IPv4 Interface Address: 70.70.70.1
Neighbor IP Address: 70.70.70.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10      IS-Extended P1.02
IPv4 Interface Address: 40.40.40.1
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec

```

```

    Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
    Metric: 10      IS-Extended P1.04
    IPv4 Interface Address: 10.10.10.2
    Neighbor IP Address: 10.10.10.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
        Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
    Metric: 10      IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
    Metric: 10      IP-Extended 70.70.70.0/24
    Prefix Attribute Flags[0]: ELC Set
    Metric: 10      IP-Extended 40.40.40.0/24
    Prefix Attribute Flags[0]: ELC Set
    Metric: 10      IP-Extended 10.10.10.0/24
    Prefix Attribute Flags[0]: ELC Set
    P1.02-00      0x00000005  0xBD05      861      0/0/0
    Metric: 0      IS-Extended P1.00
    Metric: 0      IS-Extended PE2.00
    P1.04-00      0x00000003  0x6363      603      0/0/0
    Metric: 0      IS-Extended P1.00
    Metric: 0      IS-Extended PE1.00
    P2.00-00      0x00000008  0xEE49      861      0/0/0
    Area Address: 49.0001
    NLPID: 0xCC
    Hostname: P2
    IP Address: 3.3.3.3
    Router ID: 3.3.3.3
    Metric: 10      IS-Extended P2.02
    IPv4 Interface Address: 20.20.20.2
    Neighbor IP Address: 20.20.20.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
        Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
    Metric: 10      IS-Extended P2.04
    IPv4 Interface Address: 70.70.70.2
    Neighbor IP Address: 70.70.70.2
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
        Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10

```

```

Metric: 10          IS-Extended P3.04
IPv4 Interface Address: 80.80.80.2
Neighbor IP Address: 80.80.80.1
Maximum Link Bandwidth : 1000000.00 kbits/sec
Reservable Bandwidth : 1000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P2.03
IPv4 Interface Address: 50.50.50.1
Neighbor IP Address: 50.50.50.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 3.3.3.3/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 20.20.20.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 70.70.70.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 80.80.80.0/24
Prefix Attribute Flags[0]: ELC Set
P2.02-00            0x00000005  0x6E56      644      0/0/0
Metric: 0           IS-Extended P2.00
Metric: 0           IS-Extended PE1.00
P2.03-00            0x00000005  0xB708      861      0/0/0
Metric: 0           IS-Extended P2.00
Metric: 0           IS-Extended PE2.00
P2.04-00            0x00000005  0x744D      644      0/0/0
Metric: 0           IS-Extended P2.00
Metric: 0           IS-Extended P1.00
P3.00-00            0x00000007  0x70F2      861      0/0/0
Area Address: 49.0001
NLPID: 0xCC
Hostname: P3
IP Address: 4.4.4.4
Router ID: 4.4.4.4
Metric: 10          IS-Extended P3.02
IPv4 Interface Address: 30.30.30.2
Neighbor IP Address: 30.30.30.2
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec

```



```

    Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
    Metric: 10          IS-Extended P3.04
    IPv4 Interface Address: 80.80.80.1
    Neighbor IP Address: 80.80.80.1
    Maximum Link Bandwidth : 1000000.00 kbits/sec
    Reservable Bandwidth : 1000000.00 kbits/sec
    Unreserved Bandwidth:
        Unreserved Bandwidth at priority 0: 1000000.00 kbits/sec
        Unreserved Bandwidth at priority 1: 1000000.00 kbits/sec
        Unreserved Bandwidth at priority 2: 1000000.00 kbits/sec
        Unreserved Bandwidth at priority 3: 1000000.00 kbits/sec
        Unreserved Bandwidth at priority 4: 1000000.00 kbits/sec
        Unreserved Bandwidth at priority 5: 1000000.00 kbits/sec
        Unreserved Bandwidth at priority 6: 1000000.00 kbits/sec
        Unreserved Bandwidth at priority 7: 1000000.00 kbits/sec
    TE-Default Metric: 10
    Metric: 10          IS-Extended P3.03
    IPv4 Interface Address: 60.60.60.1
    Neighbor IP Address: 60.60.60.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
        Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
    TE-Default Metric: 10
    Metric: 10          IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
    Metric: 10          IP-Extended 30.30.30.0/24
    Prefix Attribute Flags[0]: ELC Set
    Metric: 10          IP-Extended 60.60.60.0/24
    Prefix Attribute Flags[0]: ELC Set
    Metric: 10          IP-Extended 80.80.80.0/24
    Prefix Attribute Flags[0]: ELC Set
    P3.02-00          0x00000005  0x6F53          720          0/0/0
    Metric: 0          IS-Extended P3.00
    Metric: 0          IS-Extended PE1.00
    P3.03-00          0x00000005  0xB805          861          0/0/0
    Metric: 0          IS-Extended P3.00
    Metric: 0          IS-Extended PE2.00
    P3.04-00          0x00000005  0x8935          720          0/0/0
    Metric: 0          IS-Extended P3.00
    Metric: 0          IS-Extended P2.00
    PE2.00-00          0x00000006  0x839B          865          0/0/0
    Area Address: 49.0001
    NLPID: 0xCC
    Hostname: PE2
    IP Address: 5.5.5.5
    Router ID: 5.5.5.5
    Metric: 10          IS-Extended P2.03
    IPv4 Interface Address: 50.50.50.2
    Neighbor IP Address: 50.50.50.1
    Maximum Link Bandwidth : 10000000.00 kbits/sec
    Reservable Bandwidth : 10000000.00 kbits/sec
    Unreserved Bandwidth:
        Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
        Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec

```


```
Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P3.03
IPv4 Interface Address: 60.60.60.2
Neighbor IP Address: 60.60.60.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IS-Extended P1.02
IPv4 Interface Address: 40.40.40.2
Neighbor IP Address: 40.40.40.1
Maximum Link Bandwidth : 10000000.00 kbits/sec
Reservable Bandwidth : 10000000.00 kbits/sec
Unreserved Bandwidth:
  Unreserved Bandwidth at priority 0: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 1: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 2: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 3: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 4: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 5: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 6: 10000000.00 kbits/sec
  Unreserved Bandwidth at priority 7: 10000000.00 kbits/sec
TE-Default Metric: 10
Metric: 10          IP-Extended 5.5.5.5/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 50.50.50.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 60.60.60.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 40.40.40.0/24
Prefix Attribute Flags[0]: ELC Set
```

TCP MSS configuration for LDP sessions

Overview

Label Distribution Protocol (LDP) uses Transmission Control Protocol (TCP) to establish sessions between the devices. This feature enables the configuration of TCP Maximum Segment Size (MSS) that defines the maximum segment size in a single TCP segment during a communication session. TCP segment is a unit of data transmitted in a TCP connection. TCP uses three-way handshake process for initial establishment of a TCP connection. In the three-way handshake process, the sending host sends a SYN packet. Once the receiving host receives the SYN packet, it acknowledges and sends back a SYN-ACK packet to the sending host. Once the sending host receives the SYN-ACK packet from the receiving host, it sends an ACK packet, establishing a reliable connection. In this three way handshake process, the MSS is negotiated between the LDP neighbors.

Figure 10. Three-way handshake



Feature Characteristics

The configuration of the TCP MSS for LDP neighbors helps the neighbors adjust the MSS value of the TCP SYN packet. Configure the TCP MSS through the CLI and NetConf interface. The configurable MSS range is offered from 560 to 1440. By default, the MTU value for ethernet cable is 1500 bytes. When configuring the highest MSS value that is 1440, the total MSS becomes 1440 bytes (MSS) plus 20 bytes (IP Header Size), 20 bytes (TCP Header), and Ethernet header which does not cross the default path MTU value.



Note: After configuring TCP MSS, use [clear ldp session \(page 525\)](#) command to apply the MSS for the operational session.

Figure 11. Configuring TCP MSS



Benefits

By default, the interface MTU value determines the MSS value of an LDP packet. When the interface MTU value exceeds the default ethernet path MTU value of 1500 bytes, the MSS value also crosses the default ethernet path MTU value, resulting in packet fragmentation. The configuration of the specific MSS value limits the packet size irrespective of the interface MTU value, preventing packet fragmentation.[

Prerequisites

Requires the knowledge on TCP handshake and the formation of LDP neighbors.

Configuration

This section shows the procedure to configure TCP MSS for LDP session.


Enable Label Switching

Running LDP on a system requires the following tasks:

1. Enabling label-switching on the interface on NSM.
2. Enabling LDP on an interface in the LDP daemon.
3. Running an Internal Gateway Protocol (IGP), for example, Open Shortest Path first (OSPF), to distribute reachability information within the MPLS cloud.
4. Configuring the transport address.
5. Configure the TCP MSS neighbor on peer node (Active node).

Topology

Figure 12. Device topology for TCP MSS for LDP



Configuration


The below configuration shows how to configure the TCP MSS value for the LDP neighbors.

R1 - NSM

| | |
|--|--|
| R1#configure terminal | Enter configure mode. |
| R1(config)#interface ce46/1 | Specify the interface ce46/1 to be configured. |
| R1(config-if)#ip address 10.10.10.1/24 | Assign IP address 10.10.10.1/24 to interface. |
| R1(config-if)#label-switching | Enable label switching on interface ce46/1. |

| | |
|---|---|
| R1 (config-if) #exit | Exit interface mode. |
| R1 (config) #interface lo | Specify the loopback interface to be configured. |
| R1 (config-if) #ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| R1 (config-if) #commit | Commit the transaction. |

R1 - LDP

| | |
|--|---|
| R1 (config) #router ldp | Enter Router mode for LDP. |
| R1 (config-router) #router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1. |
| R1 (config-router) #transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. <div>  Note: It is preferable to use the loopback address as transport address. In addition, use the parameter <code>ipv6</code> if you are configuring an IPv6 interface. </div> |
| R1 (config-router) #targeted-peer ipv4 3.3.3.3 | Configure targeted peer 3.3.3.3. |
| R1 (config-router-targeted-peer) #exit | Exit targeted peer-mode. |
| R1 (config-router) #exit | Exit the router mode and return to the configure mode. |
| R1 (config) #interface ce46/1 | Enter interface mode ce46/1. |
| R1 (config-if) #enable-ldp ipv4 | Enable LDP on ce46/1. |
| R1 (config-if) #commit | Commit the transaction. |

R1 - OSPF

| | |
|--|---|
| R1 (config) #router ospf 100 | Configure the routing process and specify the process ID 100. The process ID should be a unique positive integer identifying the routing process. |
| R1 (config-router) #network 10.10.10.0/24 area 0 | Define the interface 10.10.10.0/24, on which OSPF runs and associate the area ID 0 with the interface. |
| R1 (config-router) #network 1.1.1.1/32 area 0 | Define the interface 1.1.1.1/32, on which OSPF runs and associate the area ID 0 with the interface. |
| R1 (config-router) #commit | Commit the transaction. |

R2 - NSM

| | |
|--|---|
| R2#configure terminal | Enter configure mode. |
| R2 (config)#interface lo | Specify the loopback interface to be configured. |
| R2 (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#interface xe0 | Specify the interface xe0 to be configured. |
| R2 (config-if)#ip address 20.20.20.1/24 | Assign IP address 20.20.20.1/24 to interface. |
| R2 (config-if)#label-switching | Enable label switching on interface xe0. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#interface xe1 | Specify the interface xe1 to be configured. |
| R2 (config-if)#ip address 10.10.10.2/24 | Assign IP address 10.10.10.2/24 to interface. |
| R2 (config-if)#label-switching | Enable label switching on interface xe1. |
| R2 (config-if)#commit | Commit the transaction. |

R2 - LDP

| | |
|---|---|
| R2 (config)#router ldp | Enter Router mode. |
| R2 (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2. |
| R2 (config-router)#transport-address ipv4 2.2.2.2 | <p>Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface.</p> <div data-bbox="868 1213 950 1276" data-label="Image"> </div> <p>Note: It is preferable to use the loopback address as transport address. In addition, use the parameter <code>ipv6</code> if you are configuring an IPv6 interface.</p> |
| R2 (config-router)#neighbor 1.1.1.1 tcp-mss 600 | Configure the TCP MSS value on peer node which have active side only. |
| R2 (config-router)#exit | Exit router mode and return to configure mode. |
| R2 (config)#interface xe0 | Specify the interface xe0 to be configured. |
| R2 (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface xe0. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#interface xe1 | Specify the interface xe1 to be configured. |
| R2 (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface xe1. |
| R2 (config-if)#commit | Commit the transaction. |

R2 - OSPF

| | |
|--|---|
| R2(config)#router ospf 100 | Configure the routing process and specify the process ID 100. The process ID should be a unique positive integer identifying the routing process. |
| R2(config-router)#network 10.10.10.0/24 area 0 | Define the interfaces 10.10.10.0/24, on which OSPF runs and associate the area ID 0 with them. |
| R2(config-router)#network 20.20.20.0/24 area 0 | Define the interfaces 20.20.20.0/24, on which OSPF runs and associate the area ID 0 with them. |
| R2(config-router)#network 2.2.2.2/32 area 0 | Define the interfaces 2.2.2.2/32, on which OSPF runs and associate the area ID 0 with them. |
| R2(config-router)#commit | Commit the transaction. |

R3 - NSM

| | |
|---|---|
| R3#configure terminal | Enter configure mode. |
| R3(config)#interface lo | Specify the loopback interface to be configured. |
| R3(config-if)#ip address 3.3.3.3/32 secondary | Set the IP address of the loopback interface to 3.3.3.3/32. |
| R3(config-if)#exit | Exit interface mode. |
| R3(config)#interface xe0 | Specify the interface xe0 to be configured. |
| R3(config-if)#ip address 20.20.20.2/24 | Set the IP address of the interface to 20.20.20.2/24. |
| R3(config-if)#label-switching | Enable label switching on interface xe0. |
| R3(config-if)#commit | Commit the transaction. |

R3 - LDP

| | |
|--|---|
| R3(config)#router ldp | Enter Router mode. |
| R3(config-router)#router-id 3.3.3.3 | Set the router ID for IP address 3.3.3.3. |
| R3(config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. <div data-bbox="873 1549 945 1612" data-label="Image"></div> Note: It is preferable to use the loopback address as transport address. In addition, use the parameter <code>ipv6</code> if you are configuring an IPv6 interface. |
| R3(config-router)#neighbor 2.2.2.2 tcp-mss 650 | Configure the TCP MSS value on peer node which have active side only. |
| R3(config-router)#targeted-peer ipv4 | Configure targeted peer. |

| | |
|--------------------------------------|--|
| 1.1.1.1 | |
| R3(config-router-targeted-peer)#exit | Exit targeted peer-mode. |
| R3(config-router)#exit | Exit the router mode and return to the configure mode. |
| R3(config)#interface xe2 | Enter interface mode xe2. |
| R3(config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| R3(config-if)#commit | Commit the transaction. |

R3 - OSPF

| | |
|--|---|
| R3(config)#router ospf 100 | Configure the routing process and specify the Process ID 100. The Process ID should be a unique positive integer identifying the routing process. |
| R3(config-router)#network 20.20.20.0/24 area 0 | Define the interfaces 20.20.20.0/24, on which OSPF runs and associate the area ID 0 with them. |
| R3(config-router)#network 3.3.3.3/32 area 0 | Define the interfaces 3.3.3.3/32, on which OSPF runs and associate the area ID 0 with them. |
| R3(config-router)#commit | Commit the transaction. |

Validation

R1

```
R1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
  2.2.2.2                ce46/1     Passive    OPERATIONAL  30        00:22:21
  3.3.3.3                ce46/1     Passive    OPERATIONAL  30        00:23:16
R1#show ldp targeted-peer count
-----
Num Targeted Peers: 1          [UP: 1]
-----
R1#show ldp session count
-----
Multicast Peers      : 1          [UP: 1]
Targeted Peers       : 1          [UP: 1]
Total Sessions       : 2          [UP: 2]
-----
R1#show ldp routes
Prefix Addr          Nexthop Addr      Intf      Owner      CreateTime  UpdateTime
1.1.1.1/32           0.0.0.0           lo        connected  00:55:55    -
2.2.2.2/32           10.10.10.2        ce46/1    ospf       00:55:34    00:17:17
3.3.3.3/32           10.10.10.2        ce46/1    ospf       00:13:42    -
10.10.10.0/24        0.0.0.0           ce46/1    connected  00:55:49    -
20.20.20.0/24        0.0.0.0           ce46/4    connected  00:55:49    -
70.70.70.0/24        10.10.10.2        ce46/1    ospf       00:16:43    -
R1#show ldp session 2.2.2.2
Session state          : OPERATIONAL
Session role           : Passive
TCP Connection         : Established
```



```

IP Address for TCP           : 2.2.2.2
Interface being used        : ce46/1
Peer LDP ID                 : 2.2.2.2:0
Last Down Reason            : Session Read Error
Preferred Peer LDP Password : Not Set
Adjacencies                 : 10.10.10.2
Advertisement mode          : Downstream Unsolicited
Label retention mode        : Liberal
Graceful Restart            : Not Capable
Keepalive Timeout           : 30
Reconnect Interval          : 15
Configured TCP MSS          : Not configured
Applied TCP MSS             : 600
Preferred TCP MSS           : NA
Address List received       : 2.2.2.2
                           : 10.10.10.2
                           : 70.70.70.1

Received Labels :      Fec      Label      Maps To
                  IPV4:3.3.3.3/32    24323      24321
                  IPV4:3.3.3.3/32    24323      24321
                  IPV4:10.10.10.0/24  impl-null  none
                  IPV4:1.1.1.1/32    24320      none
                  IPV4:70.70.70.0/24  impl-null  24322
                  IPV4:70.70.70.0/24  impl-null  24322
                  IPV4:2.2.2.2/32    impl-null  24320
                  IPV4:2.2.2.2/32    impl-null  24320

Sent Labels :      Fec      Label      Maps To
                  IPV4:3.3.3.3/32    24321      24323
                  IPV4:70.70.70.0/24  24322      impl-null
                  IPV4:20.20.20.0/24  impl-null  none
                  IPV4:10.10.10.0/24  impl-null  none
                  IPV4:2.2.2.2/32    24320      impl-null
                  IPV4:1.1.1.1/32    impl-null  none

```

```
R1#show ldp session multicast count
```

```

-----
session up count   : 1
session down count : 0
total count        : 1
-----

```

```
R1#show ldp session targeted count
```

```

-----
session up count   : 1
session down count : 0
total count        : 1
-----

```

```
R1#show ldp session all count
```

```

-----
session up count   : 2
session down count : 0
total count        : 2
-----

```

```
R1#show ldp adjacency count
```

```
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received
```

```

Type      Entity  Adjacency
-----
Multicast 1       1
Targeted  1       1
Total     2       2
-----

```

R2

```

R2#show ldp session
Codes: m - MD5 password is not set/unset.

```

```

g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 3.3.3.3 | xe0 | Passive | OPERATIONAL | 30 | 00:24:00 |
| | 1.1.1.1 | xe12 | Active | OPERATIONAL | 30 | 00:24:01 |

```
R2#show ldp targeted-peer count
```

```
-----
Num Targeted Peers: 0          [UP: 0]
-----
```

```
R2#show ldp session count
```

```
-----
Multicast Peers      : 2          [UP: 2]
Targeted Peers       : 0          [UP: 0]
Total Sessions       : 2          [UP: 2]
-----
```

```
R2#show ldp routes
```

| Prefix Addr | Nexthop Addr | Intf | Owner | CreateTime | UpdateTime |
|---------------|--------------|------|-----------|------------|------------|
| 1.1.1.1/32 | 10.10.10.1 | xe12 | ospf | 00:57:25 | 00:18:56 |
| 2.2.2.2/32 | 0.0.0.0 | lo | connected | 01:10:24 | - |
| 3.3.3.3/32 | 70.70.70.2 | xe0 | ospf | 00:15:20 | - |
| 10.10.10.0/24 | 0.0.0.0 | xe12 | connected | 00:57:28 | - |
| 70.70.70.0/24 | 0.0.0.0 | xe0 | connected | 01:09:06 | - |

```
R2#show ldp session 2.2.2.2
```

```
% No session found for address specified
```

```
R2#show ldp session multicast count
```

```
-----
session up count    : 2
session down count  : 0
total count         : 2
-----
```

```
R2#show ldp session targeted count
```

```
-----
session up count    : 0
session down count  : 0
total count         : 0
-----
```

```
R2#show ldp session all count
```

```
-----
session up count    : 2
session down count  : 0
total count         : 2
-----
```

```
R2#show ldp adjacency count
```

```
Entity - LDP enabled interfaces / configured targeted-peer
```

```
Adjacency -- Created adjacency, once hello received
```

| Type | Entity | Adjacency |
|-----------|--------|-----------|
| Multicast | 2 | 2 |
| Targeted | 0 | 0 |
| Total | 2 | 2 |

R3

```
R3#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------|-----------|--------|
|------|-----------------|---------|---------|-------|-----------|--------|

```

2.2.2.2          xe2      Active   OPERATIONAL  30    00:18:53
1.1.1.1          xe2      Active   OPERATIONAL  30    00:19:48
R3#show ldp targeted-peer count
-----
Num Targeted Peers: 1          [UP: 1]
-----
R3#show ldp session count
-----
Multicast Peers      : 1          [UP: 1]
Targeted Peers       : 1          [UP: 1]
Total Sessions       : 2          [UP: 2]
-----
R3#show ldp routes
Prefix Addr          Nexthop Addr      Intf      Owner      CreateTime  UpdateTime
1.1.1.1/32           70.70.70.1      xe2       ospf       00:10:29    -
2.2.2.2/32           70.70.70.1      xe2       ospf       00:10:29    -
3.3.3.3/32           0.0.0.0         lo        connected  01:04:19    -
10.10.10.0/24        70.70.70.1      xe2       ospf       00:10:29    -
20.20.20.0/24        0.0.0.0         xe4       connected  01:00:18    00:52:37
70.70.70.0/24        0.0.0.0         xe2       connected  01:04:15    -
R3#show ldp session 2.2.2.2
Session state        : OPERATIONAL
Session role         : Active
TCP Connection       : Established
IP Address for TCP   : 2.2.2.2
Interface being used : xe2
Peer LDP ID          : 2.2.2.2:0
Last Down Reason     : Session Read Error
Preferred Peer LDP Password : Not Set
Adjacencies          : 70.70.70.1
Advertisement mode    : Downstream Unsolicited
Label retention mode  : Liberal
Graceful Restart     : Not Capable
Keepalive Timeout    : 30
Reconnect Interval   : 15
Configured TCP MSS   : 650
Applied TCP MSS      : 650
Preferred TCP MSS     : NA
Address List received : 2.2.2.2
                      10.10.10.2
                      70.70.70.1
Received Labels :      Fec      Label      Maps To
                  IPV4:70.70.70.0/24    impl-null    none
                  IPV4:3.3.3.3/32        24323         none
                  IPV4:10.10.10.0/24     impl-null     24322
                  IPV4:10.10.10.0/24     impl-null     24322
                  IPV4:2.2.2.2/32        impl-null     24321
                  IPV4:2.2.2.2/32        impl-null     24321
                  IPV4:1.1.1.1/32        24320         24320
                  IPV4:1.1.1.1/32        24320         24320
Sent Labels :      Fec      Label      Maps To
                  IPV4:10.10.10.0/24     24322         impl-null
                  IPV4:2.2.2.2/32        24321         impl-null
                  IPV4:1.1.1.1/32        24320         24320
                  IPV4:70.70.70.0/24     impl-null     none
                  IPV4:20.20.20.0/24     impl-null     none
                  IPV4:3.3.3.3/32        impl-null     none
R3#show ldp session multicast count
-----
session up count     : 1
session down count   : 0
total count          : 1
-----
R3#show ldp session targeted count
-----
session up count     : 1
session down count   : 0
total count          : 1

```

```

-----
R3#show ldp session all count
-----
session up count    : 2
session down count  : 0
total count         : 2
-----

R3#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity  Adjacency
-----
Multicast 1      1
Targeted  1      1
Total     2      2
-----

```

Configure TCP MSS on ALL neighbor

R1 - NSM

| | |
|---|---|
| R1#configure terminal | Enter configure mode. |
| R1 (config)#interface ce46/1 | Specify the interface <code>ce46/1</code> to be configured. |
| R1 (config-if)#ip address 10.10.10.1/24 | Assign IP address 10.10.10.1/24 to interface. |
| R1 (config-if)#label-switching | Enable label switching on interface <code>ce46/1</code> . |
| R1 (config-if)#exit | Exit interface mode. |
| R1 (config)#interface lo | Specify the loopback interface to be configured. |
| R1 (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| R1 (config-if)#commit | Commit the transaction. |

R1 - LDP

| | |
|--|---|
| R1 (config)#router ldp | Enter Router mode for LDP. |
| R1 (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1. |
| R1 (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. <div data-bbox="873 1583 945 1648" data-label="Image"></div> Note: It is preferable to use the loopback address as transport address. In addition, use the parameter <code>ipv6</code> if you are configuring an IPv6 interface. |
| R1 (config-router)#targeted-peer ipv4 3.3.3.3 | Configure targeted peer. |
| R1 (config-router)#neighbor all tcp-mss | Configure the TCP MSS value with <code>all</code> neighbor. |

| | |
|--|--|
| 700 | |
| R1 (config-router-targeted-peer) #exit | Exit-targeted-peer-mode. |
| R1 (config-router) #exit | Exit the Router mode and return to the Configure mode. |
| R1 (config) #interface ce46/1 | Enter interface mode ce46/1. |
| R1 (config-if) #enable-ldp ipv4 | Enable LDP on ce46/1. |
| R1 (config-if) #commit | Commit the transaction. |

R1 - OSPF


| | |
|--|---|
| R1 (config) #router ospf 100 | Configure the routing process and specify the process ID (100). The process ID should be a unique positive integer identifying the routing process. |
| R1 (config-router) #network 10.10.10.0/24 area 0 | Define the interface 10.10.10.0/24, on which OSPF runs and associate the area ID (0) with the interface. |
| R1 (config-router) #network 1.1.1.1/32 area 0 | Define the interface 1.1.1.1/32, on which OSPF runs and associate the area ID (0) with the interface. |
| R1 (config-router) #commit | Commit the transaction. |

R2 - NSM

| | |
|---|---|
| R2#configure terminal | Enter configure mode. |
| R2 (config) #interface lo | Specify the loopback (lo) interface to be configured. |
| R2 (config-if) #ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| R2 (config-if) #exit | Exit interface mode. |
| R2 (config) #interface xe0 | Specify the interface xe0 to be configured. |
| R2 (config-if) #ip address 20.20.20.1/24 | Assign IP address 20.20.20.1/24 to interface. |
| R2 (config-if) #label-switching | Enable label switching on interface xe0. |
| R2 (config-if) #exit | Exit interface mode. |
| R2 (config) #interface xe12 | Specify the interface xe12 to be configured. |
| R2 (config-if) #ip address 10.10.10.2/24 | Assign IP address 10.10.10.2/24 to interface. |
| R2 (config-if) #label-switching | Enable label switching on interface xe12. |
| R2 (config-if) #commit | Commit the transaction. |

R2 - LDP

| | |
|-------------------------|--------------------|
| R2 (config) #router ldp | Enter Router mode. |
|-------------------------|--------------------|

| | |
|---|---|
| R2 (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2. |
| R2 (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. <div>  Note: It is preferable to use the loopback address as transport address. In addition, use the parameter <code>ipv6</code> if you are configuring an IPv6 interface. </div> |
| R2 (config-router)#neighbor all tcp-mss 710 | Configure the TCP MSS value with all neighbor. |
| R2 (config-router)#exit | Exit Router mode and return to configure mode. |
| R2 (config)#interface xe0 | Specify the interface xe0 to be configured. |
| R2 (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface xe0. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#interface xe12 | Specify the interface xe12 to be configured. |
| R2 (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface xe12. |
| R2 (config-if)#commit | Commit the transaction. |

R2 - OSPF


| | |
|---|---|
| R2 (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| R2 (config-router)#network 10.10.10.0/24 area 0 | Define the interfaces 10.10.10.0/24, on which OSPF runs and associate the area ID (0) with them. |
| R2 (config-router)#network 20.20.20.0/24 area 0 | Define the interfaces 20.20.20.0/24, on which OSPF runs and associate the area ID (0) with them. |
| R2 (config-router)#network 2.2.2.2/32 area 0 | Define the interfaces 2.2.2.2/32, on which OSPF runs and associate the area ID (0) with them. |
| R2 (config-router)#commit | Commit the transaction. |

R3 - NSM

| | |
|--|---|
| R3#configure terminal | Enter configure mode. |
| R3 (config)#interface lo | Specify the loopback interface to be configured. |
| R3 (config-if)#ip address 3.3.3.3/32 secondary | Set the IP address of the loopback interface to 3.3.3.3/32. |
| R3 (config-if)#exit | Exit interface mode. |
| R3 (config)#interface xe2 | Specify the interface xe2 to be configured. |

| | |
|--|---|
| R3(config-if)#ip address 20.20.20.2/24 | Set the IP address of the interface to 20.20.20.2/24. |
| R3(config-if)#label-switching | Enable label switching on interface2xe1. |
| R3(config-if)#commit | Commit the transaction. |

R3 - LDP

| | |
|--|---|
| R3(config)#router ldp | Enter Router mode. |
| R3(config-router)#router-id 3.3.3.3 | Set the router ID for IP address 3.3.3.3. |
| R3(config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. <div>  Note: It is preferable to use the loopback address as transport address. In addition, use the parameter <code>ipv6</code> if you are configuring an IPv6 interface. </div> |
| R3(config-router)#neighbor all tcp-mss 720 | Configure the TCP MSS value with all neighbor. |
| R3(config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| R3(config-router-targeted-peer)#exit | Exit-targeted-peer-mode. |
| R3(config-router)#exit | Exit the Router mode and return to the Configure mode. |
| R3(config)#interface xe2 | Enter interface mode. |
| R3(config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| R3(config-if)#commit | Commit the transaction. |

R3 - OSPF

| | |
|--|---|
| R3(config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| R3(config-router)#network 20.20.20.0/24 area 0 | Define the interfaces 20.20.20.0/24, on which OSPF runs and associate the area ID (0) with them. |
| R3(config-router)#network 3.3.3.3/32 area 0 | Define the interfaces 3.3.3.3/32, on which OSPF runs and associate the area ID (0) with them. |
| R3(config-router)#commit | Commit the transaction. |

Validation

R1

```
R1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
-----
      2.2.2.2              ce46/1     Passive    OPERATIONAL 30         00:01:21
      3.3.3.3              ce46/1     Passive    OPERATIONAL 30         00:01:16

R1#show ldp session 2.2.2.2
Session state           : OPERATIONAL
Session role            : Passive
TCP Connection          : Established
IP Address for TCP      : 2.2.2.2
Interface being used    : ce46/1
Peer LDP ID             : 2.2.2.2:0
Last Down Reason        : Command Line Interface (CLI) Command
Preferred Peer LDP Password : Not Set
Adjacencies             : 10.10.10.2
Advertisement mode       : Downstream Unsolicited
Label retention mode     : Liberal
Graceful Restart        : Not Capable
Keepalive Timeout       : 30
Reconnect Interval      : 15
Configured TCP MSS      : 700
Applied TCP MSS         : 600
Preferred TCP MSS        : NA
Address List received   : 2.2.2.2
                        : 10.10.10.2
                        : 70.70.70.1

Received Labels :      Fec          Label          Maps To
                  IPV4:10.10.10.0/24    impl-null      none
                  IPV4:1.1.1.1/32       24320          none
                  IPV4:70.70.70.0/24    impl-null      24322
                  IPV4:3.3.3.3/32       24321          24321
                  IPV4:2.2.2.2/32       impl-null      24320

Sent Labels :      Fec          Label          Maps To
                  IPV4:70.70.70.0/24    24322          impl-null
                  IPV4:20.20.20.0/24    impl-null      none
                  IPV4:10.10.10.0/24    impl-null      none
                  IPV4:3.3.3.3/32       24321          24321
                  IPV4:2.2.2.2/32       24320          impl-null
                  IPV4:1.1.1.1/32       impl-null      none

R1#show ldp session 3.3.3.3
Session state           : OPERATIONAL
Session role            : Passive
TCP Connection          : Established
IP Address for TCP      : 3.3.3.3
Interface being used    : ce46/1
Peer LDP ID             : 3.3.3.3:0
Last Down Reason        : Command Line Interface (CLI) Command
Preferred Peer LDP Password : Not Set
Adjacencies             : 3.3.3.3
Advertisement mode       : Downstream Unsolicited
Label retention mode     : Liberal
Graceful Restart        : Not Capable
Keepalive Timeout       : 30
Reconnect Interval      : 15
Configured TCP MSS      : 700
Applied TCP MSS         : 700
Preferred TCP MSS        : NA
Address List received   : 3.3.3.3
```



```

20.20.20.2
70.70.70.2
Received Labels :      Fec      Label      Maps To
Sent Labels :   Fec      Label      Maps To

R1#show ldp session multicast count
-----
session up count      : 1
session down count    : 0
total count           : 1
-----
R1#show ldp session targeted count
-----
session up count      : 1
session down count    : 0
total count           : 1
-----
R1#show ldp session all count
-----
session up count      : 2
session down count    : 0
total count           : 2
-----
R1#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity  Adjacency
-----
Multicast 1       1
Targeted  1       1
Total     2       2
-----

```

R2

```

R2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code Peer IP Address      IF Name    My Role    State      KeepAlive UpTime
-----
      3.3.3.3             xe0        Passive    OPERATIONAL 30      00:03:40
      1.1.1.1             xe12       Active     OPERATIONAL 30      00:02:50
R2#show ldp session 3.3.3.3
Session state           : OPERATIONAL
Session role            : Passive
TCP Connection          : Established
IP Address for TCP      : 3.3.3.3
Interface being used    : xe0
Peer LDP ID             : 3.3.3.3:0
Last Down Reason        : Command Line Interface (CLI) Command
Preferred Peer LDP Password : Not Set
Adjacencies             : 70.70.70.2
Advertisement mode       : Downstream Unsolicited
Label retention mode    : Liberal
Graceful Restart        : Not Capable
Keepalive Timeout       : 30
Reconnect Interval      : 15
Configured TCP MSS      : 710
Applied TCP MSS         : 650
Preferred TCP MSS       : NA
Address List received   : 3.3.3.3
                        20.20.20.2
                        70.70.70.2

```

```

Received Labels :      Fec          Label      Maps To
                    IPV4:70.70.70.0/24    impl-null    none
                    IPV4:20.20.20.0/24    impl-null    none
                    IPV4:10.10.10.0/24    24322       none
                    IPV4:2.2.2.2/32      24321       none
                    IPV4:1.1.1.1/32      24320       none
                    IPV4:3.3.3.3/32      impl-null    24321
                    IPV4:3.3.3.3/32      impl-null    24321
Sent Labels :      Fec          Label      Maps To
                    IPV4:70.70.70.0/24    impl-null    none
                    IPV4:10.10.10.0/24    impl-null    none
                    IPV4:3.3.3.3/32      24321       impl-null
                    IPV4:2.2.2.2/32      impl-null    none
                    IPV4:1.1.1.1/32      24320       impl-null

R2#show ldp session 1.1.1.1
Session state      : OPERATIONAL
Session role       : Active
TCP Connection     : Established
IP Address for TCP : 1.1.1.1
Interface being used : xe12
Peer LDP ID        : 1.1.1.1:0
Last Down Reason   : Session Read Error
Preferred Peer LDP Password : Not Set
Adjacencies        : 10.10.10.1
Advertisement mode  : Downstream Unsolicited
Label retention mode : Liberal
Graceful Restart   : Not Capable
Keepalive Timeout  : 30
Reconnect Interval : 15
Configured TCP MSS : 600
Applied TCP MSS    : 600
Preferred TCP MSS   : NA
Address List received : 1.1.1.1
                    10.10.10.1
                    20.20.20.1

Received Labels :      Fec          Label      Maps To
                    IPV4:70.70.70.0/24    24322       none
                    IPV4:20.20.20.0/24    impl-null    none
                    IPV4:10.10.10.0/24    impl-null    none
                    IPV4:3.3.3.3/32      24321       none
                    IPV4:2.2.2.2/32      24320       none
                    IPV4:1.1.1.1/32      impl-null    24320
                    IPV4:1.1.1.1/32      impl-null    24320
Sent Labels :      Fec          Label      Maps To
                    IPV4:70.70.70.0/24    impl-null    none
                    IPV4:10.10.10.0/24    impl-null    none
                    IPV4:3.3.3.3/32      24321       impl-null
                    IPV4:2.2.2.2/32      impl-null    none
                    IPV4:1.1.1.1/32      24320       impl-null

R2#show ldp session multicast count
-----
session up count   : 2
session down count : 0
total count        : 2
-----

R2#show ldp session targeted count
-----
session up count   : 0
session down count : 0
total count        : 0
-----

R2#show ldp session all count
-----
session up count   : 2
session down count : 0
total count        : 2
-----

R2#show ldp adjacency count

```

Entity - LDP enabled interfaces / configured targeted-peer
 Adjacency -- Created adjacency, once hello received

| Type | Entity | Adjacency |
|-----------|--------|-----------|
| ----- | | |
| Multicast | 2 | 2 |
| Targeted | 0 | 0 |
| Total | 2 | 2 |
| ----- | | |

R3

R3#show ldp session

Codes: m - MD5 password is not set/unset.
 g - GR configuration not set/unset.
 t - TCP MSS not set/unset.
 Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | xe2 | Active | OPERATIONAL | 30 | 00:04:49 |
| | 1.1.1.1 | xe2 | Active | OPERATIONAL | 30 | 00:03:55 |

R3#show ldp session 2.2.2.2

```

Session state           : OPERATIONAL
Session role            : Active
TCP Connection           : Established
IP Address for TCP       : 2.2.2.2
Interface being used     : xe2
Peer LDP ID              : 2.2.2.2:0
Last Down Reason         : Session Read Error
Preferred Peer LDP Password : Not Set
Adjacencies              : 70.70.70.1
Advertisement mode        : Downstream Unsolicited
Label retention mode     : Liberal
Graceful Restart         : Not Capable
Keepalive Timeout        : 30
Reconnect Interval       : 15
Configured TCP MSS       : 650
Applied TCP MSS          : 650
Preferred TCP MSS        : NA
Address List received    : 2.2.2.2
                          10.10.10.2
                          70.70.70.1

```

| Received Labels : | Fec | Label | Maps To |
|-------------------|--------------------|-----------|-----------|
| | IPv4:70.70.70.0/24 | impl-null | none |
| | IPv4:3.3.3.3/32 | 24321 | none |
| | IPv4:10.10.10.0/24 | impl-null | 24322 |
| | IPv4:2.2.2.2/32 | impl-null | 24321 |
| | IPv4:1.1.1.1/32 | 24320 | 24320 |
| Sent Labels : | Fec | Label | Maps To |
| | IPv4:70.70.70.0/24 | impl-null | none |
| | IPv4:20.20.20.0/24 | impl-null | none |
| | IPv4:10.10.10.0/24 | 24322 | impl-null |
| | IPv4:3.3.3.3/32 | impl-null | none |
| | IPv4:2.2.2.2/32 | 24321 | impl-null |
| | IPv4:1.1.1.1/32 | 24320 | 24320 |

R3#show ldp session 1.1.1.1

```

Session state           : OPERATIONAL
Session role            : Active
TCP Connection           : Established
IP Address for TCP       : 1.1.1.1
Interface being used     : xe2
Peer LDP ID              : 1.1.1.1:0
Last Down Reason         : Session Read Error
Preferred Peer LDP Password : Not Set
Adjacencies              : 1.1.1.1
Advertisement mode        : Downstream Unsolicited

```

```

Label retention mode      : Liberal
Graceful Restart         : Not Capable
Keepalive Timeout        : 30
Reconnect Interval       : 15
Configured TCP MSS       : 720
Applied TCP MSS          : 700
Preferred TCP MSS        : NA
Address List received     : 1.1.1.1
                          10.10.10.1
                          20.20.20.1

Received Labels :      Fec      Label      Maps To
Sent Labels :    Fec      Label      Maps To
R3#show ldp session multicast count
-----
session up count   : 1
session down count : 0
total count        : 1
-----
R3#show ldp session targeted count
-----
session up count   : 1
session down count : 0
total count        : 1
-----
R3#show ldp session all count
-----
session up count   : 2
session down count : 0
total count        : 2
-----
R3#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity  Adjacency
-----
Multicast  1       1
Targeted   1       1
Total      2       2
-----


```

Configuration of TCP MSS with Auto-targeted

R1 - NSM

| | |
|--|---|
| R1#configure terminal | Enter configure mode. |
| R1(config)#interface ce46/1 | Specify the interface ce46/1 to be configured. |
| R1(config-if)#ip address 10.10.10.1/24 | Assign IP address 10.10.10.1/24 to interface. |
| R1(config-if)#label-switching | Enable label switching on interface ce46/1. |
| R1(config-if)#exit | Exit interface mode. |
| R1(config)#interface lo | Specify the loopback interface to be configured. |
| R1(config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| R1(config-if)#commit | Commit the transaction. |

R1 - LDP

| | |
|---|---|
| R1 (config)#router ldp | Enter Router mode for LDP. |
| R1 (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1. |
| R1 (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. <div>  Note: It is preferable to use the loopback address as transport address. In addition, use the parameter <code>ipv6</code> if you are configuring an IPv6 interface. </div> |
| R1 (config-router)#targeted-peer ipv4 3.3.3.3 | Configure targeted peer. |
| R1 (config-router-targeted-peer)#exit | Exit-targeted-peer-mode. |
| R1 (config-router)#exit | Exit the Router mode and return to the configure mode. |
| R1 (config)#interface ce46/1 | Enter interface mode. |
| R1 (config-if)#enable-ldp ipv4 | Enable LDP on ce46/1. |
| R1 (config-if)#commit | Commit the transaction. |

R1 - OSPF


| | |
|---|---|
| R1 (config)#router ospf 100 | Configure the routing process and specify the process ID (100). The process ID should be a unique positive integer identifying the routing process. |
| R1 (config-router)#network 10.10.10.0/24 area 0 | Define the interface 10.10.10.0/24, on which OSPF runs and associate the area ID (0) with the interface. |
| R1 (config-router)#network 1.1.1.1/32 area 0 | Define the interface 1.1.1.1/32, on which OSPF runs and associate the area ID (0) with the interface. |
| R1 (config-router)#commit | Commit the transaction. |

R2 - NSM

| | |
|--|---|
| R2#configure terminal | Enter configure mode. |
| R2 (config)#interface lo | Specify the loopback interface to be configured. |
| R2 (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#interface xe0 | Specify the interface xe0 to be configured. |

| | |
|---|---|
| R2 (config-if)#ip address 20.20.20.1/24 | Assign IP address 20.20.20.1/24 to interface. |
| R2 (config-if)#label-switching | Enable label switching on interface xe0. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#interface xe12 | Specify the interface xe12 to be configured. |
| R2 (config-if)#ip address 10.10.10.2/24 | Assign IP address 10.10.10.2/24 to interface. |
| R2 (config-if)#label-switching | Enable label switching on interface xe12. |
| R2 (config-if)#commit | Commit the transaction. |

R2 - LDP

| | |
|---|--|
| R2 (config)#router ldp | Enter Router mode. |
| R2 (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2. |
| R2 (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. <div>  Note: It is preferable to use the loopback address as transport address. In addition, use the parameter <code>ipv6</code> if you are configuring an IPv6 interface. </div> |
| R2 (config-router)#neighbor auto-targeted tcp-mss 800 | Configure the TCP MSS value on all auto-targeted neighbors. |
| R2 (config-router)#exit | Exit Router mode and return to configure mode. |
| R2 (config)#interface xe0 | Specify the interface xe0 to be configured. |
| R2 (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface xe0. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#interface xe12 | Specify the interface xe12 to be configured. |
| R2 (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface xe12. |
| R2 (config-if)#commit | Commit the transaction. |

R2 - OSPF


| | |
|---|---|
| R2 (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| R2 (config-router)#network 10.10.10.0/24 area 0 | Define the interfaces 10.10.10.0/24, on which OSPF runs and associate the area ID (0) with them. |
| R2 (config-router)#network 20.20.20.0/24 area 0 | Define the interfaces 20.20.20.0/24, on which OSPF runs and associate the area ID (0) with them. |

| | |
|--|---|
| R2 (config-router)#network 2.2.2.2/32 area 0 | Define the interfaces 2.2.2.2/32, on which OSPF runs and associate the area ID (0) with them. |
| R2 (config-router)#commit | Commit the transaction. |

R3 - NSM

| | |
|--|---|
| R3#configure terminal | Enter configure mode. |
| R3 (config)#interface lo | Specify the loopback interface to be configured. |
| R3 (config-if)#ip address 3.3.3.3/32 secondary | Set the IP address of the loopback interface to 3.3.3.3/32. |
| R3 (config-if)#exit | Exit interface mode. |
| R3 (config)#interface xe2 | Specify the interface xe2 to be configured. |
| R3 (config-if)#ip address 20.20.20.2/24 | Set the IP address of the interface to 20.20.20.2/24. |
| R3 (config-if)#label-switching | Enable label switching on interface xe2. |
| R3 (config-if)#commit | Commit the transaction. |

R3 - LDP

| | |
|---|--|
| R3 (config)#router ldp | Enter Router mode. |
| R3 (config-router)#router-id 3.3.3.3 | Set the router ID for IP address 3.3.3.3. |
| R3 (config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. <div>  Note: It is preferable to use the loopback address as transport address. In addition, use the parameter <code>ipv6</code> if you are configuring an IPv6 interface. </div> |
| R3 (config-router)#neighbor auto-targeted tcp-mss 810 | Configure the TCP MSS value on all auto-targeted neighbors. |
| R3 (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| R3 (config-router-targeted-peer)#exit | Exit-targeted-peer-mode. |
| R3 (config-router)#exit | Exit the Router mode and return to the configure mode. |
| R3 (config)#interface xe2 | Enter interface mode xe2. |
| R3 (config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| R3 (config-if)#commit | Commit the transaction. |

R3 - OSPF

| | |
|--|---|
| R3(config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| R3(config-router)#network 20.20.20.0/24 area 0 | Define the interfaces 20.20.20.0/24, on which OSPF runs and associate the area ID (0) with them. |
| R3(config-router)#network 3.3.3.3/32 area 0 | Define the interfaces 3.3.3.3/32, on which OSPF runs and associate the area ID (0) with them. |
| R3(config-router)#commit | Commit the transaction. |

Validation**R1**

```

R1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
     2.2.2.2              ce46/1     Passive    OPERATIONAL 30      00:01:41
     3.3.3.3              ce46/1     Passive    OPERATIONAL 30      00:01:32

R1#
R1#sh ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
     2.2.2.2              ce46/1     Passive    OPERATIONAL 30      00:02:00
     3.3.3.3              ce46/1     Passive    OPERATIONAL 30      00:01:51

R1#show ldp targeted-peers
IP Address      Interface
3.3.3.3         ce46/1

R1#sh ldp session 2.2.2.2
Session state           : OPERATIONAL
Session role            : Passive
TCP Connection          : Established
IP Address for TCP      : 2.2.2.2
Interface being used    : ce46/1
Peer LDP ID             : 2.2.2.2:0
Last Down Reason        : Command Line Interface (CLI) Command
Preferred Peer LDP Password : Not Set
Adjacencies             : 10.10.10.2
Advertisement mode       : Downstream Unsolicited
Label retention mode     : Liberal
Graceful Restart        : Not Capable
Keepalive Timeout       : 30
Reconnect Interval      : 15
Configured TCP MSS      : 700
Applied TCP MSS         : 600
Preferred TCP MSS        : NA
Address List received    : 2.2.2.2
                        10.10.10.2
                        70.70.70.1

Received Labels :      Fec              Label              Maps To

```



```

                IPV4:10.10.10.0/24      impl-null    none
                IPV4:1.1.1.1/32         24320         none
                IPV4:70.70.70.0/24      impl-null     24322
                IPV4:3.3.3.3/32         24321         24321
                IPV4:2.2.2.2/32         impl-null     24320
Sent Labels :   Fec                    Label         Maps To
                IPV4:70.70.70.0/24      24322         impl-null
                IPV4:20.20.20.0/24      impl-null     none
                IPV4:10.10.10.0/24      impl-null     none
                IPV4:3.3.3.3/32         24321         24321
                IPV4:2.2.2.2/32         24320         impl-null
                IPV4:1.1.1.1/32         impl-null     none

```

R1#sh ldp session 3.3.3.3

```

Session state      : OPERATIONAL
Session role       : Passive
TCP Connection     : Established
IP Address for TCP : 3.3.3.3
Interface being used : ce46/1
Peer LDP ID        : 3.3.3.3:0
Last Down Reason   : Session Read Error
Preferred Peer LDP Password : Not Set
Adjacencies        : 3.3.3.3
Advertisement mode  : Downstream Unsolicited
Label retention mode : Liberal
Graceful Restart   : Not Capable
Keepalive Timeout  : 30
Reconnect Interval : 15
Configured TCP MSS : 700
Applied TCP MSS    : 700
Preferred TCP MSS   : NA
Address List received : 3.3.3.3
                   20.20.20.2
                   70.70.70.2

```

```

Received Labels :      Fec          Label      Maps To
Sent Labels :      Fec          Label      Maps To

```

R1#show ldp session multicast count

```

-----
session up count   : 1
session down count : 0
total count        : 1
-----

```

R1#show ldp session targeted count

```

-----
session up count   : 1
session down count : 0
total count        : 1
-----

```

R1#show ldp session all count

```

-----
session up count   : 2
session down count : 0
total count        : 2
-----

```

R1#show ldp adjacency count

Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

```

Type      Entity  Adjacency
-----
Multicast 1       1
Targeted  1       1
Total     2       2
-----

```

R2

```

R2#show ldp session
Codes: m - MD5 password is not set/unset.
      g - GR configuration not set/unset.
      t - TCP MSS not set/unset.
      Session has to be cleared manually

Code Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
     3.3.3.3              xe0        Passive    OPERATIONAL 30         00:04:17
     1.1.1.1              xe12       Active     OPERATIONAL 30         00:04:29

R2#show ldp targeted-peer
targeted-peer    targeted-peers

R2#show ldp targeted-peers
R2#show ldp session 3.3.3.3
Session state      : OPERATIONAL
Session role       : Passive
TCP Connection     : Established
IP Address for TCP : 3.3.3.3
Interface being used : xe0
Peer LDP ID        : 3.3.3.3:0
Last Down Reason   : Command Line Interface (CLI) Command
Preferred Peer LDP Password : Not Set
Adjacencies        : 70.70.70.2
Advertisement mode  : Downstream Unsolicited
Label retention mode : Liberal
Graceful Restart   : Not Capable
Keepalive Timeout  : 30
Reconnect Interval : 15
Configured TCP MSS : 710
Applied TCP MSS    : 650
Preferred TCP MSS   : NA
Address List received : 3.3.3.3
                   20.20.20.2
                   70.70.70.2

Received Labels :      Fec      Label      Maps To
                IPV4:70.70.70.0/24    impl-null    none
                IPV4:20.20.20.0/24    impl-null    none
                IPV4:10.10.10.0/24    24322        none
                IPV4:2.2.2.2/32       24321        none
                IPV4:1.1.1.1/32       24320        none
                IPV4:3.3.3.3/32       impl-null    24321
                IPV4:3.3.3.3/32       impl-null    24321

Sent Labels :      Fec      Label      Maps To
                IPV4:70.70.70.0/24    impl-null    none
                IPV4:10.10.10.0/24    impl-null    none
                IPV4:3.3.3.3/32       24321        impl-null
                IPV4:2.2.2.2/32       impl-null    none
                IPV4:1.1.1.1/32       24320        impl-null

R2#show ldp session 1.1.1.1
Session state      : OPERATIONAL
Session role       : Active
TCP Connection     : Established
IP Address for TCP : 1.1.1.1
Interface being used : xe12
Peer LDP ID        : 1.1.1.1:0
Last Down Reason   : Session Read Error
Preferred Peer LDP Password : Not Set
Adjacencies        : 10.10.10.1
Advertisement mode  : Downstream Unsolicited
Label retention mode : Liberal
Graceful Restart   : Not Capable
Keepalive Timeout  : 30
Reconnect Interval : 15
Configured TCP MSS : 600

```

```

Applied TCP MSS      : 600
Preferred TCP MSS    : NA
Address List received : 1.1.1.1
                     10.10.10.1
                     20.20.20.1

Received Labels :      Fec      Label      Maps To
                  IPV4:70.70.70.0/24  24322      none
                  IPV4:20.20.20.0/24  impl-null   none
                  IPV4:10.10.10.0/24  impl-null   none
                  IPV4:3.3.3.3/32     24321      none
                  IPV4:2.2.2.2/32     24320      none
                  IPV4:1.1.1.1/32     impl-null   24320
                  IPV4:1.1.1.1/32     impl-null   24320

Sent Labels :      Fec      Label      Maps To
                  IPV4:70.70.70.0/24  impl-null   none
                  IPV4:10.10.10.0/24  impl-null   none
                  IPV4:3.3.3.3/32     24321      impl-null
                  IPV4:2.2.2.2/32     impl-null   none
                  IPV4:1.1.1.1/32     24320      impl-null

R2#show ldp session multicast count
-----
session up count   : 2
session down count : 0
total count        : 2
-----

R2#show ldp session targeted count
-----
session up count   : 0
session down count : 0
total count        : 0
-----

R2#show ldp session all count
-----
session up count   : 2
session down count : 0
total count        : 2
-----

R2#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity  Adjacency
-----
Multicast 2       2
Targeted  0       0
Total     2       2
-----

```

R3

```

R3#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name  My Role  State      KeepAlive  UpTime
     2.2.2.2              xe2     Active  OPERATIONAL  30      00:06:20
     1.1.1.1              xe2     Active  OPERATIONAL  30      00:06:20

R3#show ldp targeted-peers
IP Address      Interface
1.1.1.1         xe2

R3#show ldp session 2.2.2.2
Session state    : OPERATIONAL
Session role     : Active
TCP Connection   : Established

```

```

IP Address for TCP           : 2.2.2.2
Interface being used        : xe2
Peer LDP ID                 : 2.2.2.2:0
Last Down Reason            : Session Read Error
Preferred Peer LDP Password : Not Set
Adjacencies                 : 70.70.70.1
Advertisement mode           : Downstream Unsolicited
Label retention mode        : Liberal
Graceful Restart            : Not Capable
Keepalive Timeout           : 30
Reconnect Interval          : 15
Configured TCP MSS          : 650
Applied TCP MSS             : 650
Preferred TCP MSS           : NA
Address List received       : 2.2.2.2
                           : 10.10.10.2
                           : 70.70.70.1

Received Labels :      Fec      Label      Maps To
                  IPV4:70.70.70.0/24    impl-null    none
                  IPV4:3.3.3.3/32        24321        none
                  IPV4:10.10.10.0/24     impl-null    24322
                  IPV4:2.2.2.2/32        impl-null    24321
                  IPV4:1.1.1.1/32        24320        24320

Sent Labels :      Fec      Label      Maps To
                  IPV4:70.70.70.0/24    impl-null    none
                  IPV4:20.20.20.0/24    impl-null    none
                  IPV4:10.10.10.0/24    24322        impl-null
                  IPV4:3.3.3.3/32        impl-null    none
                  IPV4:2.2.2.2/32        24321        impl-null
                  IPV4:1.1.1.1/32        24320        24320

R3#show ldp session 1.1.1.1
Session state           : OPERATIONAL
Session role            : Active
TCP Connection          : Established
IP Address for TCP      : 1.1.1.1
Interface being used    : xe2
Peer LDP ID             : 1.1.1.1:0
Last Down Reason        : Command Line Interface (CLI) Command
Preferred Peer LDP Password : Not Set
Adjacencies             : 1.1.1.1
Advertisement mode       : Downstream Unsolicited
Label retention mode     : Liberal
Graceful Restart         : Not Capable
Keepalive Timeout       : 30
Reconnect Interval      : 15
Configured TCP MSS      : 720
Applied TCP MSS          : 700
Preferred TCP MSS        : NA
Address List received    : 1.1.1.1
                        : 10.10.10.1
                        : 20.20.20.1

Received Labels :      Fec      Label      Maps To
Sent Labels :      Fec      Label      Maps To
R3#show ldp session multicast count
-----
session up count   : 1
session down count : 0
total count        : 1
-----
R3#show ldp session targeted count
-----
session up count   : 1
session down count : 0
total count        : 1
-----
R3#show ldp session all count
-----
session up count   : 2

```

```

session down count : 0
total count       : 2
-----
R3#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity  Adjacency
-----
Multicast  1       1
Targeted   1       1
Total      2       2
-----

```

TCP MSS for LDP Session Command

neighbor tcp-mss

Use this command to set the TCP MSS for an LDP session. MSS is a TCP parameter that defines the maximum amount of data in a TCP segment that can be transmitted.

Use the `no` command to remove the TCP MSS from an LDP session.

Command Syntax

```

neighbor (A.B.C.D | auto-targeted | all) tcp-mss <560-1440>
no neighbor (A.B.C.D | auto-targeted | all) tcp-mss

```

Parameters

A.B.C.D

To set MSS for the specific peer.

auto-targeted

To set MSS for auto-targeted LDP peer. Auto-targeted LDP sessions automatically establish the TCP connection with neighboring routers and do not require the manual configuration of each peer.

all

To set MSS for all LDP peers

<560-1440>

Configure the TCP MSS between this range. The default value is 1460 bytes.

Default

Disabled

Command Mode

Router LDP mode.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS(config)#router ldp
```

```
OcNOS(config-router)#neighbor 2.2.2.2 tcp-mss 900
OcNOS(config-router)#neighbor all tcp-mss 1000
OcNOS(config-router)#neighbor auto-targeted tcp-mss 800
OcNOS(config-router)#commit
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

| Acronym | Description |
|---------|-------------------------------|
| ACK | Acknowledgment |
| IGP | Interior Gateway Protocol |
| LDP | Label Distribution Protocol |
| MSS | Maximum Segment Size |
| MTU | Maximum Transmission Unit |
| OSPF | Open Short Path First |
| SYN | Synchronize |
| TCP | Transmission Control Protocol |

Glossary

The following provides definitions for key terms used throughout this document:

| | |
|-------------|---|
| LDP | LDP is a routing protocol that manages and distributes the labels to the route in a Multiprotocol Label Switching (MPLS) network. Adding a label to a route helps to control the flow of network traffic and increases the forwarding speed, ensuring a smooth and optimized data transmission. |
| LDP session | LDP session is the connection established between LDP routers in an MPLS network. |
| MSS | MSS is a TCP parameter that defines the maximum amount of data in a TCP segment that can be transmitted. |
| TCP | TCP is one of the main protocols in the Internet Protocol (IP) suite. It offers a secure and reliable connection between two devices. |
| TCP segment | TCP segment is a unit of data transmitted in a TCP connection. The segment consists of header and payload. The header contains the control information to manage the transmission, and the payload contains the actual data that needs to be transmitted. |

LDP Remote Loop-Free Alternate (RLFA)

A basic mechanism using Loop-Free Alternates (LFAs) is described in RFC5286 that provides good repair coverage in many topologies, especially those that are highly meshed.

However, some topologies, notably ring-based topologies, are not well protected by LFAs alone. This is because there is no neighbor of the Point of Local Repair (PLR) that has a cost to the destination via a path that does not traverse the failure that is cheaper than the cost to the destination via the failure.

RFC 7490 provides extensions to the basic repair mechanism in which tunnels are used to provide additional logical links that can be used as loop-free alternates where none exist in the original topology. It provides loop-free alternates that guarantee only link protection.

RFC 8102 provides remote-loop-free-based IP fast reroute mechanisms that specifies procedures for determining whether or not a given PQ-node provides node protection for a specific destination. It provides node protection for all destinations covered by the same remote-LFA alternate, in case of failure of the primary next-hop node


ISIS shall compute PQ node and LDP shall dynamically create tunnel to PQ node so that if primary path fails traffic can be rerouted to backup rLFA tunnel and hence to destination.

Remote LFA involves the use of a tunnel to a next-hop that is not directly connected. This is the primary difference between the LFA and Remote LFA.

RLFA Terminology

Terms used for defining tunnel as below:

Figure 13. Understanding RLFA



Repair tunnel

A tunnel established for the purpose of providing a virtual neighbor that is a Loop-Free Alternate.

P-space

The P-space of a router with respect to a protected link is the set of routers reachable from that specific router using the pre-convergence shortest paths, without any of those paths (including equal cost path splits) transiting that protected link.

For example, the P-space of S with respect to link S-E is the set of routers that S can reach without using the protected link S-E.

Extended P-space

Consider the set of neighbors of a router protecting a link. Exclude the router reachable over the protected link from that set of routers. The extended P-space of the protecting router with respect to the protected link is the union of the P spaces of the neighbors in that set of neighbors with respect to the protected link.

Q-space

Q-space of a router with respect to a protected link is the set of routers from which that specific router that can be reached without any path (including ECMP Splits) transiting that protected link.

PQ node

A PQ node of a node S with respect to a protected link S-E is a node that is a member of both the P-space (or the extended P-space) of S with respect to that protected link S-E and the Q-space of E with respect to that protected link S-E. A repair tunnel endpoint is chosen from the set of PQ-nodes.

Remote LFA (RLFA)

The use of a PQ node rather than a neighbor of the repairing node as the next hop in an LFA repair.

In [Figure 13](#), S can reach A, B, and C without going via S-E; these form S's extended P-space with respect to S-E. The routers that can reach E without going through S-E will be in E's Q-space with respect to link S-E; these are D and C. B has equal-cost paths to E via B-A-S-E and B-C-D-E, and so the forwarder at S might choose to send a packet to E via link S-E. Hence, B is not in the Q-space of E with respect to link S-E. The single node in both S's extended P-space and E's Q-space is C; thus, node C is selected as the repair tunnel's endpoint. Thus, if a tunnel is provided between S and C as shown in [Figure 14. RLFA Topology \(page 382\)](#), then C, now being a direct neighbor of S, would become an LFA for D and E.

Establishing RLFA Tunnel

To calculate the Remote LFA backup path and to determine the Remote LFA node the software requirement can be broadly classified as below:

IS-IS

ISIS shall calculate Repair path that are P space (routers it can reach without traversing the protected link) and Q space (routers that can reach the protected destination without traversing the protected link). Hence routers that

belong to both spaces called as PQ routers has to be calculated. It has to inform LDP about PQ node so that RLFA tunnel can be established.

LDP

LDP shall establish targeted session with PQ node and shall advertise label to peer node for the destination FEC. LDP shall send primary and backup path FTN/ILM add/delete info to NSM for further programing.

NSM

NSM shall maintain Primary and Backup path FTN/ILM and shall send information to HSL for data-plane programing.

HSL

HSL program primary and backup FTN/ILM entry in hardware.

Configure LDP Remote Loop-Free Alternate (RLFA)

The LDP RLFA configuration process can be divided into the following tasks:

1. Enable label-switching on the interface on NSM.
2. Establish IS-IS routing between the nodes (to distribute reachability information within the MPLS cloud)
3. Configure IS-IS RLFA on Source node
4. Enabling LDP on an interface in the LDP daemon
5. Configure LDP FRR with Auto-targeted-session (Allow creating TLDP session dynamically)
6. Enable BFD interval globally and for all ISIS enabled interfaces



Notes:

- Faster convergence can be achieved with lower BFD interval enabled globally.
- Dynamically created RLFA T-LDP sessions will be removed only after disabling LDP `auto-targeted-session` CLI or LDP FRR.
- When `targeted-peer ipv4` CLI is configured with `auto-targeted-session` CLI, T-LDP session created for targeted-peer only remove after disabling `auto-targeted-session` CLI.
- After Enabling ISIS RLFA, Both LFA and RLFA computation will be done and RLFA path will be preferred to provide node-protection.
- Better convergence can be achieved with LDP-IGP-SYNC enabled.

Assumptions and Limitations


- Only LDP(MPLS) serves as the tunnel mechanism to reach a Remote-LFA repair node.
- Only IPv4 protocol is supported.
- RFC 7916 (LFA-MANAGE) is not supported.
- RLFA next-hop does not support ECMP.

- LDP does not support RLFA backup paths that are resolved over another LFA or RLFA backup path. Consequently, if the IGP computes an RLFA backup for a given prefix, LDP will not install that path if it requires a next-hop that itself relies on a non-primary LFA or RLFA path.

Topology

The below topology shows the configuration required to enable the RLFA feature.

Figure 14. RLFA Topology



PE1 Configurations

| | |
|---------------------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To Allow creating TLDP session dynamically |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit LDP process |
| (config)#interface ce6/2 | Enter interface mode. |
| (config-if)#ip address 10.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa) |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on ce6/2 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |

| | |
|---|--|
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe29</code> | Enter interface mode. |
| <code>(config-if)#ip address 60.1.1.2/24</code> | Configure the IP address of the interface. |
| <code>(config-if)#ip router isis rlfa</code> | Enable IS-IS routing on an interface for area 49 (rlfa). |
| <code>(config-if)#label-switching</code> | Enable label-switching on interface |
| <code>(config-if)#enable-ldp ipv4</code> | Enable ldp process on xe29 interface |
| <code>(config-if)#commit</code> | Commit the candidate configuration to the running configuration |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface lo</code> | Enter interface mode. |
| <code>(config-if)#ip address 1.1.1.1/32 secondary</code> | Configure the IP address of the interface |
| <code>(config-if)#ip router isis rlfa</code> | Enable IS-IS routing on an interface for area 49 (rlfa). |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router isis rlfa</code> | Create an IS-IS routing instance for area 49 (rlfa). |
| <code>(config-router)#is-type level-1</code> | Configure instance as level-1-only routing. |
| <code>(config-router)#metric-style wide</code> | Configure the new style of metric type as wide. |
| <code>(config-router)#mpls traffic-eng level-1</code> | Enable MPLS-TE in is-type Level-1. |
| <code>(config-router)# mpls traffic-eng router-id 1.1.1.1</code> | Configure MPLS-TE unique router-id TLV. |
| <code>(config-router)#dynamic-hostname</code> | Configure the hostname to be advertised for an ISIS instance. |
| <code>(config-router)# net 49.0000.0000.0001.00</code> | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| <code>(config-router)# fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp</code> | Configure Remote LFA to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| <code>(config-router)# bfd all-interfaces</code> | Enable the Bidirectional Forwarding Detection (BFD) feature on the interfaces enabled with this ISIS instance. |
| <code>(config-router)#commit</code> | Commit the candidate configuration to the running configuration |
| <code>(config-router)#end</code> | Exit router mode. |
| <code>(config)# bfd interval 3 minrx 3 multiplier 3</code> | Configure bfd interval globally |

PE2 Configurations

| | |
|--|---|
| <code>#configure terminal</code> | Enter configure mode. |
| <code>(config)#router ldp</code> | Enable LDP process |
| <code>(config-router)#fast-reroute</code> | Enable LDP FRR |
| <code>(config-router)#auto-targeted-session</code> | To Allow creating TLDP session dynamically |
| <code>(config-router)#commit</code> | Commit the candidate configuration to the running |

| | |
|---|--|
| | configuration |
| (config-router)#exit | Exit LDP process |
| (config)#interface xe17 | Enter interface mode. |
| (config-if)#ip address 20.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa) |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on xe17 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface cd0/2 | Enter interface mode. |
| (config-if)#ip address 30.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on cd0/2 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure the IP address of the interface |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis rlfa | Create an IS-IS routing instance for area 49 (rlfa). |
| (config-router)#is-type level-1 | Configure instance as level-1-only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1. |
| (config-router)# mpls traffic-eng router-id 3.3.3.3 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| (config-router)# net 49.0000.0000.0003.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)# bfd all-interfaces | Enable the Bidirectional Forwarding Detection (BFD) feature on the interfaces enabled with this ISIS instance. |
| (config-router)#commit | Commit the candidate configuration to the running configuration |

| | |
|---|---|
| (config-router)#exit | Exit router mode. |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |
| (config)#commit | Commit the candidate configuration to the running configuration |

PE3 Configurations

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To Allow creating TLDP session dynamically |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit LDP process |
| (config)#interface xe23 | Enter interface mode. |
| (config-if)#ip address 30.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa) |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on xe23 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 40.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on xe2 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure the IP address of the interface |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis rlfa | Create an IS-IS routing instance for area 49 (rlfa). |
| (config-router)#is-type level-1 | Configure instance as level-1-only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1. |

| | |
|---|--|
| (config-router)# mpls traffic-eng router-id 4.4.4.4 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| (config-router)# net 49.0000.0000.0004.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)# bfd all-interfaces | Enable the Bidirectional Forwarding Detection (BFD) feature on the interfaces enabled with this ISIS instance. |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit router mode. |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |
| (config)#commit | Commit the candidate configuration to the running configuration |

PE4 Configurations

| | |
|---------------------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To Allow creating TLDP session dynamically |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit LDP process |
| (config)#interface xe6 | Enter interface mode. |
| (config-if)#ip address 50.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa) |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on xe6 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 60.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on xe2 interface |

| | |
|---|--|
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 6.6.6.6/32 secondary | Configure the IP address of the interface |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis rlfa | Create an IS-IS routing instance for area 49 (rlfa). |
| (config-router)#is-type level-1 | Configure instance as level-1-only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1. |
| (config-router)# mpls traffic-eng router-id 6.6.6.6 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| (config-router)# net 49.0000.0000.0005.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)# bfd all-interfaces | Enable the Bidirectional Forwarding Detection (BFD) feature on the interfaces enabled with this ISIS instance. |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit router mode. |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |
| (config)#commit | Commit the candidate configuration to the running configuration |

P1 Configurations

| | |
|---------------------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To Allow creating TLDP session dynamically |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit LDP process |
| (config)#interface xe8 | Enter interface mode. |
| (config-if)#ip address 10.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa) |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on xe8 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |

| | |
|---|--|
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe17 | Enter interface mode. |
| (config-if)#ip address 20.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on xe17 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure the IP address of the interface |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis rlfa | Create an IS-IS routing instance for area 49 (rlfa). |
| (config-router)#is-type level-1 | Configure instance as level-1-only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1. |
| (config-router)# mpls traffic-eng router-id 1.1.1.1 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| (config-router)# net 49.0000.0000.0002.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)# bfd all-interfaces | Enable the Bidirectional Forwarding Detection (BFD) feature on the interfaces enabled with this ISIS instance. |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit router mode. |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |
| (config)#commit | Commit the candidate configuration to the running configuration |

P2 Configurations

| | |
|------------------------------|-----------------------|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |

| | |
|---|--|
| (config-router)#auto-targeted-session | To Allow creating TLDP session dynamically |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit LDP process |
| (config)#interface ge4 | Enter interface mode. |
| (config-if)#ip address 40.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa) |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on ge4 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface ge6 | Enter interface mode. |
| (config-if)#ip address 50.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#label-switching | Enable label-switching on interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on ge6 interface |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 5.5.5.5/32 secondary | Configure the IP address of the interface |
| (config-if)#ip router isis rlfa | Enable IS-IS routing on an interface for area 49 (rlfa). |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis rlfa | Create an IS-IS routing instance for area 49 (rlfa). |
| (config-router)#is-type level-1 | Configure instance as level-1-only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1. |
| (config-router)# mpls traffic-eng router-id 5.5.5.5 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| (config-router)# net 49.0000.0000.0005.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)# bfd all-interfaces | Enable the Bidirectional Forwarding Detection (BFD) feature on the interfaces enabled with this ISIS instance. |
| (config-router)#commit | Commit the candidate configuration to the running configuration |

| | |
|---|---|
| (config-router)#exit | Exit router mode. |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |
| (config)#commit | Commit the candidate configuration to the running configuration |

Validation

Verify LDP Neighborhood Before Enabling RLFA

PE1

```
PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | ce6/2 | Passive | OPERATIONAL | 30 | 00:53:26 |
| | 6.6.6.6 | xe29 | Passive | OPERATIONAL | 30 | 00:53:25 |

P1

```
P1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | xe8 | Active | OPERATIONAL | 30 | 00:54:07 |
| | 3.3.3.3 | xe17 | Passive | OPERATIONAL | 30 | 00:54:05 |

PE2

```
PE2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | xe17 | Active | OPERATIONAL | 30 | 00:55:46 |
| | 4.4.4.4 | cd0/2 | Passive | OPERATIONAL | 30 | 00:55:46 |

PE3

```
PE3#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 3.3.3.3 | xe23 | Active | OPERATIONAL | 30 | 00:56:19 |
| | 5.5.5.5 | xe2 | Passive | OPERATIONAL | 30 | 00:56:19 |
| | 1.1.1.1 | xe23 | Active | OPERATIONAL | 30 | 00:56:17 |

P2

```
P2#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 4.4.4.4 | ge4 | Active | OPERATIONAL | 30 | 00:56:28 |
| | 6.6.6.6 | ge6 | Passive | OPERATIONAL | 30 | 00:56:28 |

PE4

```
PE4#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 5.5.5.5 | xe6 | Active | OPERATIONAL | 30 | 00:56:43 |
| | 1.1.1.1 | xe2 | Active | OPERATIONAL | 30 | 00:56:41 |

Verify IS-IS Adjacency

Check the output of "show clns neighbors" to confirm that the IS-IS adjacency is active (up).

PE1

```
PE1#show clns neighbors
```

```
Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
```

```
Tag rlfa: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| PE4 | xe29 | b86a.97f5.ebc8 | Up | 27 | L1 | IS-IS |
| P1 | ce6/2 | 80a2.355b.680a | Up | 21 | L1 | IS-IS |

P1

```
P1#show clns neighbors
```

```
Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
```

```
Tag rlfa: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| PE1 | xe8 | e001.a6aa.1120 | Up | 4 | L1 | IS-IS |
| PE2 | xe17 | e8c5.7ada.68e1 | Up | 6 | L1 | IS-IS |

PE2

```
PE2#show clns neighbors
```

```
Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
```

```
Tag rlfa: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| PE3 | cd0/2 | b86a.97f5.e0dd | Up | 21 | L1 | IS-IS |
| P1 | xe17 | 80a2.355b.6813 | Up | 28 | L1 | IS-IS |

PE3

```
PE3#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag rlfa: VRF : default
System Id      Interface  SNPA              State Holdtime  Type Protocol
P2              xe2        e8c5.7a69.4bce    Up    7          L1    IS-IS
PE2             xe23       e8c5.7ada.68ec    Up    6          L1    IS-IS
```

P2

```
P2#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag rlfa: VRF : default
System Id      Interface  SNPA              State Holdtime  Type Protocol
PE3             ge4        b86a.97f5.e0c8    Up    28         L1    IS-IS
PE4             ge6        b86a.97f5.ebcc    Up    27         L1    IS-IS
```

PE4

```
PE4#show clns neighbors

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag rlfa: VRF : default
System Id      Interface  SNPA              State Holdtime  Type Protocol
PE1             xe2        e001.a6aa.111f    Up    6          L1    IS-IS
P2              xe6        e8c5.7a69.4bd0    Up    6          L1    IS-IS
```

Verify IS-IS Route Installation

Check the IS-IS route installation in the IS-IS table and RIB table.

PE1

```
PE1#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag rlfa: VRF : default
Total number of routes: 12

  Destination      Metric      Next-Hop      Interface      Tag
C    1.1.1.1/32      10          --            lo              0
L1   2.2.2.2/32      20          10.1.1.2      ce6/2           0
L1   3.3.3.3/32      30          10.1.1.2      ce6/2           0
L1   4.4.4.4/32      40          60.1.1.1      xe29            0
      10.1.1.2      ce6/2           0
L1   5.5.5.5/32      30          60.1.1.1      xe29            0
L1   6.6.6.6/32      20          60.1.1.1      xe29            0
C    10.1.1.0/24      10          --            ce6/2           0
L1   20.1.1.0/24      20          10.1.1.2      ce6/2           0
L1   30.1.1.0/24      30          10.1.1.2      ce6/2           0
L1   40.1.1.0/24      30          60.1.1.1      xe29            0
```

```

L1  50.1.1.0/24      20      60.1.1.1      xe29      0
C   60.1.1.0/24      10      --            xe29      0

```

```
PE1#show ip route
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
```

```
O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
```

```
ia - IS-IS inter area, E - EVPN,
```

```
v - vrf leaked
```

```
* - candidate default
```

```
IP Route Table for VRF "default"
```

```

C       1.1.1.1/32 is directly connected, lo, installed 04:50:49, last update 04:50:49 ago
i L1    2.2.2.2/32 [115/20] via 10.1.1.2, ce6/2, installed 04:43:29, last update 04:43:29 ago
i L1    3.3.3.3/32 [115/30] via 10.1.1.2, ce6/2, installed 04:22:48, last update 04:22:48 ago
i L1    4.4.4.4/32 [115/40] via 10.1.1.2, ce6/2, installed 04:19:04, last update 01:20:03 ago
         [115/40] via 60.1.1.1, xe29
i L1    5.5.5.5/32 [115/30] via 60.1.1.1, xe29, installed 01:20:03, last update 01:20:03 ago
i L1    6.6.6.6/32 [115/20] via 60.1.1.1, xe29, installed 01:20:03, last update 01:20:03 ago
C       10.1.1.0/24 is directly connected, ce6/2, installed 04:51:41, last update 04:51:41 ago
i L1    20.1.1.0/24 [115/20] via 10.1.1.2, ce6/2, installed 04:43:29, last update 04:43:29 ago
i L1    30.1.1.0/24 [115/30] via 10.1.1.2, ce6/2, installed 04:22:48, last update 04:22:48 ago
i L1    40.1.1.0/24 [115/30] via 60.1.1.1, xe29, installed 01:20:03, last update 01:20:03 ago
i L1    50.1.1.0/24 [115/20] via 60.1.1.1, xe29, installed 01:20:03, last update 01:20:03 ago
C       60.1.1.0/24 is directly connected, xe29, installed 04:51:12, last update 04:51:12 ago
C       127.0.0.0/8 is directly connected, lo, installed 04:57:44, last update 04:57:44 ago

```

```
Gateway of last resort is not set
```

P1

```
P1#show ip isis route
```

```
Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
```

```
ia - IS-IS inter area, D - discard, e - external metric
```

```
** - invalid
```

```
Tag rlfa: VRF : default
```

```
Total number of routes: 12
```

| | Destination | Metric | Next-Hop | Interface | Tag |
|----|-------------|--------|----------|-----------|-----|
| L1 | 1.1.1.1/32 | 20 | 10.1.1.1 | xe8 | 0 |
| C | 2.2.2.2/32 | 10 | -- | lo | 0 |
| L1 | 3.3.3.3/32 | 20 | 20.1.1.2 | xe17 | 0 |
| L1 | 4.4.4.4/32 | 30 | 20.1.1.2 | xe17 | 0 |
| L1 | 5.5.5.5/32 | 40 | 20.1.1.2 | xe17 | 0 |
| | | | 10.1.1.1 | xe8 | 0 |
| L1 | 6.6.6.6/32 | 30 | 10.1.1.1 | xe8 | 0 |
| C | 10.1.1.0/24 | 10 | -- | xe8 | 0 |
| C | 20.1.1.0/24 | 10 | -- | xe17 | 0 |
| L1 | 30.1.1.0/24 | 20 | 20.1.1.2 | xe17 | 0 |
| L1 | 40.1.1.0/24 | 30 | 20.1.1.2 | xe17 | 0 |
| L1 | 50.1.1.0/24 | 30 | 10.1.1.1 | xe8 | 0 |
| L1 | 60.1.1.0/24 | 20 | 10.1.1.1 | xe8 | 0 |

```
P1#show ip route
```

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
```

```
O - OSPF, IA - OSPF inter area
```

```
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
```

```
E1 - OSPF external type 1, E2 - OSPF external type 2
```

```
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
```

```
ia - IS-IS inter area, E - EVPN,
```

```

v - vrf leaked
* - candidate default

IP Route Table for VRF "default"
i L1      1.1.1.1/32 [115/20] via 10.1.1.1, xe8, installed 04:44:06, last update 04:44:06 ago
C         2.2.2.2/32 is directly connected, lo, installed 04:44:59, last update 04:44:59 ago
i L1      3.3.3.3/32 [115/20] via 20.1.1.2, xe17, installed 04:23:22, last update 04:23:22 ago
i L1      4.4.4.4/32 [115/30] via 20.1.1.2, xe17, installed 04:19:39, last update 04:19:39 ago
i L1      5.5.5.5/32 [115/40] via 10.1.1.1, xe8, installed 04:16:38, last update 01:20:38 ago
          [115/40] via 20.1.1.2, xe17
i L1      6.6.6.6/32 [115/30] via 10.1.1.1, xe8, installed 01:20:38, last update 01:20:38 ago
C         10.1.1.0/24 is directly connected, xe8, installed 04:45:48, last update 04:45:48 ago
C         20.1.1.0/24 is directly connected, xe17, installed 04:45:21, last update 04:45:21 ago
i L1      30.1.1.0/24 [115/20] via 20.1.1.2, xe17, installed 04:23:22, last update 04:23:22 ago
i L1      40.1.1.0/24 [115/30] via 20.1.1.2, xe17, installed 04:19:39, last update 04:19:39 ago
i L1      50.1.1.0/24 [115/30] via 10.1.1.1, xe8, installed 01:20:38, last update 01:20:38 ago
i L1      60.1.1.0/24 [115/20] via 10.1.1.1, xe8, installed 04:44:06, last update 04:44:06 ago
C         127.0.0.0/8 is directly connected, lo, installed 04:55:25, last update 04:55:25 ago

Gateway of last resort is not set

```

PE2

```

PE2#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

Tag rifa: VRF : default
Total number of routes: 12

      Destination      Metric      Next-Hop      Interface      Tag
L1  1.1.1.1/32         30      20.1.1.1      xe17            0
L1  2.2.2.2/32         20      20.1.1.1      xe17            0
C   3.3.3.3/32         10      --           lo              0
L1  4.4.4.4/32         20      30.1.1.2      cd0/2           0
L1  5.5.5.5/32         30      30.1.1.2      cd0/2           0
L1  6.6.6.6/32         40      30.1.1.2      cd0/2           0
          20.1.1.1      xe17            0
L1  10.1.1.0/24        20      20.1.1.1      xe17            0
C   20.1.1.0/24        10      --           xe17            0
C   30.1.1.0/24        10      --           cd0/2           0
L1  40.1.1.0/24        20      30.1.1.2      cd0/2           0
L1  50.1.1.0/24        30      30.1.1.2      cd0/2           0
L1  60.1.1.0/24        30      20.1.1.1      xe17            0

```

```

PE2#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
i L1      1.1.1.1/32 [115/30] via 20.1.1.1, xe17, installed 04:24:22, last update 04:24:22 ago
i L1      2.2.2.2/32 [115/20] via 20.1.1.1, xe17, installed 04:24:34, last update 04:24:34 ago
C         3.3.3.3/32 is directly connected, lo, installed 04:25:22, last update 04:25:22 ago
i L1      4.4.4.4/32 [115/20] via 30.1.1.2, cd0/2, installed 04:20:48, last update 04:20:48 ago
i L1      5.5.5.5/32 [115/30] via 30.1.1.2, cd0/2, installed 04:17:48, last update 04:17:48 ago
i L1      6.6.6.6/32 [115/40] via 20.1.1.1, xe17, installed 01:21:47, last update 01:21:47 ago
          [115/40] via 30.1.1.2, cd0/2
i L1      10.1.1.0/24 [115/20] via 20.1.1.1, xe17, installed 04:24:34, last update 04:24:34 ago
C         20.1.1.0/24 is directly connected, xe17, installed 04:26:43, last update 04:26:43 ago

```

```

C          30.1.1.0/24 is directly connected, cd0/2, installed 04:25:38, last update 04:25:38 ago
i L1      40.1.1.0/24 [115/20] via 30.1.1.2, cd0/2, installed 04:20:48, last update 04:20:48 ago
i L1      50.1.1.0/24 [115/30] via 30.1.1.2, cd0/2, installed 04:17:48, last update 04:17:48 ago
i L1      60.1.1.0/24 [115/30] via 20.1.1.1, xe17, installed 04:24:22, last update 04:24:22 ago
C          127.0.0.0/8 is directly connected, lo, installed 04:59:31, last update 04:59:31 ago

```

Gateway of last resort is not set

PE3

```
PE3#show ip isis route
```

```

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, D - discard, e - external metric
       ** - invalid

```

```

Tag r1fa: VRF : default
Total number of routes: 12

```

| | Destination | Metric | Next-Hop | Interface | Tag |
|----|-------------|--------|----------|-----------|-----|
| L1 | 1.1.1.1/32 | 40 | 40.1.1.2 | xe2 | 0 |
| | | | 30.1.1.1 | xe23 | 0 |
| L1 | 2.2.2.2/32 | 30 | 30.1.1.1 | xe23 | 0 |
| L1 | 3.3.3.3/32 | 20 | 30.1.1.1 | xe23 | 0 |
| C | 4.4.4.4/32 | 10 | -- | lo | 0 |
| L1 | 5.5.5.5/32 | 20 | 40.1.1.2 | xe2 | 0 |
| L1 | 6.6.6.6/32 | 30 | 40.1.1.2 | xe2 | 0 |
| L1 | 10.1.1.0/24 | 30 | 30.1.1.1 | xe23 | 0 |
| L1 | 20.1.1.0/24 | 20 | 30.1.1.1 | xe23 | 0 |
| C | 30.1.1.0/24 | 10 | -- | xe23 | 0 |
| C | 40.1.1.0/24 | 10 | -- | xe2 | 0 |
| L1 | 50.1.1.0/24 | 20 | 40.1.1.2 | xe2 | 0 |
| L1 | 60.1.1.0/24 | 30 | 40.1.1.2 | xe2 | 0 |

```
PE3#show ip route
```

```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

```

```
IP Route Table for VRF "default"
```

```

i L1      1.1.1.1/32 [115/40] via 30.1.1.1, xe23, installed 04:21:22, last update 01:22:32 ago
           [115/40] via 40.1.1.2, xe2
i L1      2.2.2.2/32 [115/30] via 30.1.1.1, xe23, installed 04:21:22, last update 04:21:22 ago
i L1      3.3.3.3/32 [115/20] via 30.1.1.1, xe23, installed 04:21:35, last update 04:21:35 ago
C          4.4.4.4/32 is directly connected, lo, installed 04:22:51, last update 04:22:51 ago
i L1      5.5.5.5/32 [115/20] via 40.1.1.2, xe2, installed 04:18:32, last update 04:18:32 ago
i L1      6.6.6.6/32 [115/30] via 40.1.1.2, xe2, installed 01:22:32, last update 01:22:32 ago
i L1      10.1.1.0/24 [115/30] via 30.1.1.1, xe23, installed 04:21:22, last update 04:21:22 ago
i L1      20.1.1.0/24 [115/20] via 30.1.1.1, xe23, installed 04:21:35, last update 04:21:35 ago
C          30.1.1.0/24 is directly connected, xe23, installed 04:23:50, last update 04:23:50 ago
C          40.1.1.0/24 is directly connected, xe2, installed 04:23:12, last update 04:23:12 ago
i L1      50.1.1.0/24 [115/20] via 40.1.1.2, xe2, installed 04:18:32, last update 04:18:32 ago
i L1      60.1.1.0/24 [115/30] via 40.1.1.2, xe2, installed 01:22:32, last update 01:22:32 ago
C          127.0.0.0/8 is directly connected, lo, installed 04:59:49, last update 04:59:49 ago

```

Gateway of last resort is not set

P2

```
P2#show ip isis route
```

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
 ia - IS-IS inter area, D - discard, e - external metric
 ** - invalid

Tag r1fa: VRF : default
 Total number of routes: 12

| | Destination | Metric | Next-Hop | Interface | Tag |
|----|-------------|--------|----------|-----------|-----|
| L1 | 1.1.1.1/32 | 30 | 50.1.1.2 | ge6 | 0 |
| L1 | 2.2.2.2/32 | 40 | 50.1.1.2 | ge6 | 0 |
| | | | 40.1.1.1 | ge4 | 0 |
| L1 | 3.3.3.3/32 | 30 | 40.1.1.1 | ge4 | 0 |
| L1 | 4.4.4.4/32 | 20 | 40.1.1.1 | ge4 | 0 |
| C | 5.5.5.5/32 | 10 | -- | lo | 0 |
| L1 | 6.6.6.6/32 | 20 | 50.1.1.2 | ge6 | 0 |
| L1 | 10.1.1.0/24 | 30 | 50.1.1.2 | ge6 | 0 |
| L1 | 20.1.1.0/24 | 30 | 40.1.1.1 | ge4 | 0 |
| L1 | 30.1.1.0/24 | 20 | 40.1.1.1 | ge4 | 0 |
| C | 40.1.1.0/24 | 10 | -- | ge4 | 0 |
| C | 50.1.1.0/24 | 10 | -- | ge6 | 0 |
| L1 | 60.1.1.0/24 | 20 | 50.1.1.2 | ge6 | 0 |

P2#show ip route

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
 O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
 ia - IS-IS inter area, E - EVPN,
 v - vrf leaked
 * - candidate default

IP Route Table for VRF "default"

| | | | | | |
|------|-------------|----------|-----------------------------|---------------------|--------------------------|
| i L1 | 1.1.1.1/32 | [115/30] | via 50.1.1.2, ge6, | installed 01:23:25, | last update 01:23:25 ago |
| i L1 | 2.2.2.2/32 | [115/40] | via 40.1.1.1, ge4, | installed 04:19:15, | last update 01:23:25 ago |
| | | | [115/40] via 50.1.1.2, ge6 | | |
| i L1 | 3.3.3.3/32 | [115/30] | via 40.1.1.1, ge4, | installed 04:19:15, | last update 04:19:15 ago |
| i L1 | 4.4.4.4/32 | [115/20] | via 40.1.1.1, ge4, | installed 04:19:27, | last update 04:19:27 ago |
| C | 5.5.5.5/32 | | is directly connected, lo, | installed 04:20:10, | last update 04:20:10 ago |
| i L1 | 6.6.6.6/32 | [115/20] | via 50.1.1.2, ge6, | installed 01:23:25, | last update 01:23:25 ago |
| i L1 | 10.1.1.0/24 | [115/30] | via 50.1.1.2, ge6, | installed 01:23:25, | last update 01:23:25 ago |
| i L1 | 20.1.1.0/24 | [115/30] | via 40.1.1.1, ge4, | installed 04:19:15, | last update 04:19:15 ago |
| i L1 | 30.1.1.0/24 | [115/20] | via 40.1.1.1, ge4, | installed 04:19:27, | last update 04:19:27 ago |
| C | 40.1.1.0/24 | | is directly connected, ge4, | installed 04:21:05, | last update 04:21:05 ago |
| C | 50.1.1.0/24 | | is directly connected, ge6, | installed 04:20:26, | last update 04:20:26 ago |
| i L1 | 60.1.1.0/24 | [115/20] | via 50.1.1.2, ge6, | installed 01:23:25, | last update 01:23:25 ago |
| C | 127.0.0.0/8 | | is directly connected, lo, | installed 05:01:40, | last update 05:01:40 ago |

Gateway of last resort is not set

PE4

PE4#show ip isis route

Codes: C - connected, E - external, L1 - IS-IS level-1, L2 - IS-IS level-2
 ia - IS-IS inter area, D - discard, e - external metric
 ** - invalid

Tag r1fa: VRF : default
 Total number of routes: 12

| | Destination | Metric | Next-Hop | Interface | Tag |
|----|-------------|--------|----------|-----------|-----|
| L1 | 1.1.1.1/32 | 20 | 60.1.1.2 | xe2 | 0 |
| L1 | 2.2.2.2/32 | 30 | 60.1.1.2 | xe2 | 0 |
| L1 | 3.3.3.3/32 | 40 | 50.1.1.1 | xe6 | 0 |
| | | | 60.1.1.2 | xe2 | 0 |


```

L1 4.4.4.4/32      30      50.1.1.1      xe6      0
L1 5.5.5.5/32      20      50.1.1.1      xe6      0
C 6.6.6.6/32      10      --           lo       0
L1 10.1.1.0/24     20      60.1.1.2      xe2      0
L1 20.1.1.0/24     30      60.1.1.2      xe2      0
L1 30.1.1.0/24     30      50.1.1.1      xe6      0
L1 40.1.1.0/24     20      50.1.1.1      xe6      0
C 50.1.1.0/24     10      --           xe6      0
C 60.1.1.0/24     10      --           xe2      0

```

```
PE4#show ip route
```

```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

```

```
IP Route Table for VRF "default"
```

```

i L1      1.1.1.1/32 [115/20] via 60.1.1.2, xe2, installed 01:24:33, last update 01:24:33 ago
i L1      2.2.2.2/32 [115/30] via 60.1.1.2, xe2, installed 01:24:21, last update 01:24:21 ago
i L1      3.3.3.3/32 [115/40] via 60.1.1.2, xe2, installed 01:24:21, last update 01:24:21 ago
           [115/40] via 50.1.1.1, xe6
i L1      4.4.4.4/32 [115/30] via 50.1.1.1, xe6, installed 01:24:21, last update 01:24:21 ago
i L1      5.5.5.5/32 [115/20] via 50.1.1.1, xe6, installed 01:24:33, last update 01:24:33 ago
C         6.6.6.6/32 is directly connected, lo, installed 01:25:32, last update 01:25:32 ago
i L1      10.1.1.0/24 [115/20] via 60.1.1.2, xe2, installed 01:24:33, last update 01:24:33 ago
i L1      20.1.1.0/24 [115/30] via 60.1.1.2, xe2, installed 01:24:21, last update 01:24:21 ago
i L1      30.1.1.0/24 [115/30] via 50.1.1.1, xe6, installed 01:24:21, last update 01:24:21 ago
i L1      40.1.1.0/24 [115/20] via 50.1.1.1, xe6, installed 01:24:33, last update 01:24:33 ago
C         50.1.1.0/24 is directly connected, xe6, installed 01:26:24, last update 01:26:24 ago
C         60.1.1.0/24 is directly connected, xe2, installed 01:25:54, last update 01:25:54 ago
C         127.0.0.0/8 is directly connected, lo, installed 05:01:50, last update 05:01:50 ago

```

```
Gateway of last resort is not set
```

Verify Primary Path Computation

Verify the computed paths for Primary Paths in the IS-IS LFA and RLFA backup.

PE1

```
PE1#show ip isis route fast-reroute
```

```

Tag   : rlfa  VRF : default
Total number of routes: 9

```

```

Codes : L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area,
       D - discard, LP - Link Protecting, NP - Node Protecting,
       BP - Broadcast Interface Disjoint, Pri - Primary Path,
       Sec - Secondary Path, DP - Downstream Path

```

```

L1 2.2.2.2/32
  Primary Path via      : 10.1.1.2, ce6/2
  Remote FRR Path via   : 4.4.4.4, via : 60.1.1.1, xe29
  FRR Metric            : 40
  Protection Provided   : LP

```

```

L1 3.3.3.3/32
  Primary Path via      : 10.1.1.2, ce6/2
  Remote FRR Path via   : 4.4.4.4, via : 60.1.1.1, xe29
  FRR Metric            : 40

```

```

Protection Provided : LP

L1 4.4.4.4/32
Primary Path via      : 60.1.1.1, xe29
FRR Backup Path via   : 10.1.1.2, ce6/2
FRR Metric             : 40
Protection Provided    : LP NP BP Pri DP

Primary Path via      : 10.1.1.2, ce6/2
FRR Backup Path via   : 60.1.1.1, xe29
FRR Metric             : 40
Protection Provided    : LP NP BP Pri DP

L1 5.5.5.5/32
Primary Path via      : 60.1.1.1, xe29
Remote FRR Path via   : 4.4.4.4, via : 10.1.1.2, ce6/2
FRR Metric             : 40
Protection Provided    : LP

L1 6.6.6.6/32
Primary Path via      : 60.1.1.1, xe29
Remote FRR Path via   : 4.4.4.4, via : 10.1.1.2, ce6/2
FRR Metric             : 40
Protection Provided    : LP

L1 20.1.1.0/24
Primary Path via      : 10.1.1.2, ce6/2
Remote FRR Path via   : 4.4.4.4, via : 60.1.1.1, xe29
FRR Metric             : 40
Protection Provided    : LP

L1 30.1.1.0/24
Primary Path via      : 10.1.1.2, ce6/2
FRR Backup Path via   : 60.1.1.1, xe29
FRR Metric             : 40
Protection Provided    : LP NP BP

L1 40.1.1.0/24
Primary Path via      : 60.1.1.1, xe29
FRR Backup Path via   : 10.1.1.2, ce6/2
FRR Metric             : 40
Protection Provided    : LP NP BP

L1 50.1.1.0/24
Primary Path via      : 60.1.1.1, xe29
Remote FRR Path via   : 4.4.4.4, via : 10.1.1.2, ce6/2
FRR Metric             : 40
Protection Provided    : LP

```

PQ Node Verification

Verify that the PQ node near the source is selected, and the Target-LDP session is established with the PQ node using the commands below.

```

PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
-----
2.2.2.2      ce6/2      Passive    OPERATIONAL  30        01:05:50
6.6.6.6      xe29       Passive    OPERATIONAL  30        01:05:49
4.4.4.4      ce6/2      Passive    OPERATIONAL  30        01:05:49

```

```

PE1#show ldp targeted-peers
IP Address      Interface
4.4.4.4         ce6/2  > PE1 established T-LDP with PE3 (since PE3 is PQ for PE1)

```

FTN Installation Verification

Verify that the Primary and Backup FTNs are installed with labels in the LDP RLFA route table.

```

PE1#show ldp rlfa-routes
Codes: p - stale rLFA route
Fec          Primary-NH      Backup-NH      rLFA-Addr      Out-Intf  Outer-label  Inner-label
Owner
2.2.2.2      10.1.1.2      60.1.1.1      4.4.4.4        xe29      24323      24321      isi
s
3.3.3.3      10.1.1.2      60.1.1.1      4.4.4.4        xe29      24323      24322      isi
s
5.5.5.5      60.1.1.1      10.1.1.2      4.4.4.4        ce6/2     24322      24323      isi
s
6.6.6.6      60.1.1.1      10.1.1.2      4.4.4.4        ce6/2     24322      24324      isi
s
20.1.1.0     10.1.1.2      60.1.1.1      4.4.4.4        xe29      24323      24326      isi
s
50.1.1.0     60.1.1.1      10.1.1.2      4.4.4.4        ce6/2     24322      24327      isi
s

```

Backup Cross Connect (XC) Verification in FTN Table

Verify that the backup XCs are calculated for the primary FTNs in the MPLS forwarding table and confirm the same in the FTN table.

PE1

```

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code  FEC          FTN-ID  Nhlf-ID  Tunnel-ID  Pri  Out-Label  Out-
Intf  ELC          Nexthop  Algo-Num  UpTime
L>    2.2.2.2/32      1      18      -      -      -      -
      -              N/A      01:06:54
      1              -
      Yes  3              ce6/2    No      10.1.1.2      -      -
      No  24321      xe29     No      4.4.4.4      -      -
      60.1.1.1, label 24323)
      L>  3.3.3.3/32      2      39      -      -      -      -
      -              N/A      01:06:54
      10             -
      Yes  24321      ce6/2    No      10.1.1.2      -      -
      No  24322      xe29     No      4.4.4.4      -      -
      60.1.1.1, label 24323)
      L>  4.4.4.4/32      3      16      -      -      -      -
      -              N/A      01:06:54
      3              -
      Yes  24322      ce6/2    No      10.1.1.2      -      -
      No  24323      xe29     No      60.1.1.1      -      -

```

```

        Yes    24323      xe29      15      -      60.1.1.1      -      -
        No    24322      ce6/2      3      -      10.1.1.2      -      -
L> 5.5.5.5/32      6      61      -      -      -      -
        -      N/A      01:06:54      31      -
        Yes    24324      xe29      60      -      60.1.1.1      -      -
        No    24323      ce6/2      60      -      4.4.4.4      -      -
10.1.1.2, label 24322)
L> 6.6.6.6/32      7      65      -      -      -      -
        -      N/A      01:06:54      36      -
        Yes    3      xe29      64      -      60.1.1.1      -      -
        No    24324      ce6/2      64      -      4.4.4.4      -      -
10.1.1.2, label 24322)
L> 20.1.1.0/24      4      51      -      -      -      -
        -      N/A      01:06:54      1      -
        Yes    3      ce6/2      47      -      10.1.1.2      -      -
        No    24326      xe29      47      -      4.4.4.4      -      -
60.1.1.1, label 24323)
L> 30.1.1.0/24      5      69      -      -      -      -
        -      N/A      01:06:54      57      -
        Yes    24325      ce6/2      68      -      10.1.1.2      -      -
        No    24327      xe29      72      -      60.1.1.1      -      -
L> 40.1.1.0/24      8      72      -      -      -      -
        -      N/A      01:06:54      44      -
        Yes    24328      xe29      71      -      60.1.1.1      -      -
        No    24326      ce6/2      76      -      10.1.1.2      -      -
L> 50.1.1.0/24      9      76      -      -      -      -
        -      N/A      01:06:54      36      -
        Yes    3      xe29      75      -      60.1.1.1      -      -
        No    24327      ce6/2      75      -      4.4.4.4      -      -
10.1.1.2, label 24322)

PE1#show mpls ftn-table
Primary FTM entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:05, UpTime: 01:07:05, LastUpdate: 01:07:05
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 8, out intf: ce6/2, out label: 3
Nexthop addr: 10.1.1.2      cross connect ix: 1, op code: Push

Backup Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 17 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 17, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24321
Nexthop addr: 4.4.4.4      cross connect ix: 3, op code: Push

Primary FTM entry with FEC: 3.3.3.3/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed

```

```
CreateTime: 01:07:05, UpTime: 01:07:05, LastUpdate: 01:07:05
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 10 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 10, owner: LDP, Stale: NO, refcount: 2, out intf: ce6/2, out label: 24321
  Nexthop addr: 10.1.1.2      cross connect ix: 4, op code: Push

  Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 34 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 34, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24322
  Nexthop addr: 4.4.4.4      cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:05, UpTime: 01:07:05, LastUpdate: 01:07:05
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 3 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 3, owner: LDP, Stale: NO, refcount: 6, out intf: ce6/2, out label: 24322
  Nexthop addr: 10.1.1.2      cross connect ix: 2, op code: Push

  Backup Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 15 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 15, owner: LDP, Stale: NO, refcount: 4, out intf: xe29, out label: 24323
  Nexthop addr: 60.1.1.1      cross connect ix: 2, op code: Push

  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 15 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 15, owner: LDP, Stale: NO, refcount: 4, out intf: xe29, out label: 24323
  Nexthop addr: 60.1.1.1      cross connect ix: 2, op code: Push

  Backup Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 3 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 3, owner: LDP, Stale: NO, refcount: 6, out intf: ce6/2, out label: 24322
  Nexthop addr: 10.1.1.2      cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:05, UpTime: 01:07:05, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 31 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 31, owner: LDP, Stale: NO, refcount: 3, out intf: xe29, out label: 24324
  Nexthop addr: 60.1.1.1      cross connect ix: 5, op code: Push

  Backup Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 60 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 60, owner: LDP, Stale: NO, refcount: 2, out intf: ce6/2, out label: 24323
  Nexthop addr: 4.4.4.4      cross connect ix: 11, op code: Push

Primary FTN entry with FEC: 6.6.6.6/32, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:05, UpTime: 01:07:05, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 36 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 36, owner: N/A, Stale: NO, refcount: 7, out intf: xe29, out label: 3
```

```
Nexthop addr: 60.1.1.1      cross connect ix: 7, op code: Push

Backup Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 64 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 64, owner: LDP, Stale: NO, refcount: 2, out intf: ce6/2, out label: 24324
Nexthop addr: 4.4.4.4      cross connect ix: 13, op code: Push

Primary FTN entry with FEC: 20.1.1.0/24, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:05, UpTime: 01:07:05, LastUpdate: 01:07:05
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 8, out intf: ce6/2, out label: 3
Nexthop addr: 10.1.1.2      cross connect ix: 1, op code: Push

Backup Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 47 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 47, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24326
Nexthop addr: 4.4.4.4      cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 30.1.1.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:05, UpTime: 01:07:05, LastUpdate: 01:07:05
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 57 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 57, owner: LDP, Stale: NO, refcount: 2, out intf: ce6/2, out label: 24325
Nexthop addr: 10.1.1.2      cross connect ix: 9, op code: Push

Backup Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 68 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 68, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24327
Nexthop addr: 60.1.1.1      cross connect ix: 15, op code: Push

Primary FTN entry with FEC: 40.1.1.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:05, UpTime: 01:07:05, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 44 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 44, owner: LDP, Stale: NO, refcount: 3, out intf: xe29, out label: 24328
Nexthop addr: 60.1.1.1      cross connect ix: 8, op code: Push

Backup Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 71 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 71, owner: LDP, Stale: NO, refcount: 1, out intf: ce6/2, out label: 24326
Nexthop addr: 10.1.1.2      cross connect ix: 17, op code: Push

Primary FTN entry with FEC: 50.1.1.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:05, UpTime: 01:07:05, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 36 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 36, owner: N/A, Stale: NO, refcount: 7, out intf: xe29, out label: 3
Nexthop addr: 60.1.1.1      cross connect ix: 7, op code: Push
```

```

Backup Cross connect ix: 18, in intf: - in label: 0 out-segment ix: 75 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 75, owner: LDP, Stale: NO, refcount: 2, out intf: ce6/2, out label: 24327
Nexthop addr: 4.4.4.4 cross connect ix: 18, op code: Push

```

P1

```
Pl#show mpls forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | Nexthop | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|-------------|---------|----------|----------|-----------|-----|-----------|------|
| Intf | ELC | | Algo-Num | UpTime | | | | |
| L> | 1.1.1.1/32 | | 1 | 15 | - | - | - | - |
| | - | N/A | | 01:07:41 | | | | |
| | Yes 3 | | xe8 | No | 10.1.1.1 | - | - | - |
| L> | 3.3.3.3/32 | | 6 | 32 | - | - | - | - |
| | - | N/A | | 01:07:41 | | | | |
| | Yes 3 | | xe17 | No | 20.1.1.2 | - | - | - |
| L> | 4.4.4.4/32 | | 7 | 35 | - | - | - | - |
| | - | N/A | | 01:07:41 | | | | |
| | Yes 24322 | | xe17 | No | 20.1.1.2 | - | - | - |
| L> | 5.5.5.5/32 | | 2 | 38 | - | - | - | - |
| | - | N/A | | 01:07:41 | | | | |
| | Yes 24323 | | xe17 | No | 20.1.1.2 | - | - | - |
| | | | | 17 | - | | | |
| L> | 6.6.6.6/32 | | 3 | 27 | - | - | - | - |
| | - | N/A | | 01:07:41 | | | | |
| | Yes 24324 | | xe8 | No | 10.1.1.1 | - | - | - |
| L> | 30.1.1.0/24 | | 8 | 32 | - | - | - | - |
| | - | N/A | | 01:07:41 | | | | |
| | Yes 3 | | xe17 | No | 20.1.1.2 | - | - | - |
| L> | 40.1.1.0/24 | | 9 | 39 | - | - | - | - |
| | - | N/A | | 01:07:41 | | | | |
| | Yes 24326 | | xe17 | No | 20.1.1.2 | - | - | - |
| L> | 50.1.1.0/24 | | 4 | 30 | - | - | - | - |
| | - | N/A | | 01:07:41 | | | | |
| | Yes 24328 | | xe8 | No | 10.1.1.1 | - | - | - |
| L> | 60.1.1.0/24 | | 5 | 15 | - | - | - | - |
| | - | N/A | | 01:07:41 | | | | |
| | Yes 3 | | xe8 | No | 10.1.1.1 | - | - | - |

```
Pl#show mpls ftn-table
```

```
Primary FTN entry with FEC: 1.1.1.1/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed
```

```
CreateTime: 01:07:44, UpTime: 01:07:44, LastUpdate: N/A
```

```
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number: 0
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1
```

```
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 5, out intf: xe8, out label: 3
```

Nexthop addr: 10.1.1.1 cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 6, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:07:44, UpTime: 01:07:44, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 5, out intf: xe17, out label: 3

Nexthop addr: 20.1.1.2 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 7, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:07:44, UpTime: 01:07:44, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 34 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 34, owner: LDP, Stale: NO, refcount: 2, out intf: xe17, out label: 24322

Nexthop addr: 20.1.1.2 cross connect ix: 8, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 2, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:07:44, UpTime: 01:07:44, LastUpdate: 01:07:44

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 37 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 37, owner: LDP, Stale: NO, refcount: 1, out intf: xe17, out label: 24323

Nexthop addr: 20.1.1.2 cross connect ix: 9, op code: Push

Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 17 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 17, owner: LDP, Stale: NO, refcount: 1, out intf: xe8, out label: 24323

Nexthop addr: 10.1.1.1 cross connect ix: 9, op code: Push

Primary FTN entry with FEC: 6.6.6.6/32, id: 3, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:07:44, UpTime: 01:07:44, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 26 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 26, owner: LDP, Stale: NO, refcount: 2, out intf: xe8, out label: 24324

Nexthop addr: 10.1.1.1 cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 30.1.1.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:07:44, UpTime: 01:07:44, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 5, out intf: xe17, out label: 3

Nexthop addr: 20.1.1.2 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 40.1.1.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State: Installed


```

CreateTime: 01:07:44, UpTime: 01:07:44, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 13 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 13, owner: LDP, Stale: NO, refcount: 3, out intf: xe17, out label: 24326
Nexthop addr: 20.1.1.2      cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 50.1.1.0/24, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:44, UpTime: 01:07:44, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 29 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 29, owner: LDP, Stale: NO, refcount: 2, out intf: xe8, out label: 24328
Nexthop addr: 10.1.1.1      cross connect ix: 7, op code: Push

Primary FTN entry with FEC: 60.1.1.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:07:44, UpTime: 01:07:44, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 5, out intf: xe8, out label: 3
Nexthop addr: 10.1.1.1      cross connect ix: 1, op code: Push

```

PE2

```

PE2#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code   FEC           FTN-ID   Nhlfe-ID  Tunnel-ID  Pri   Out-Label  Out-
Intf   ELC           Nexthop   Algo-Num  UpTime
L>     1.1.1.1/32    1         23        -          -     -          -
      -          N/A         01:08:26
      Yes 24320    xe17      No        20.1.1.1  -     -          -
L>     2.2.2.2/32    7         37        -          -     -          -
      -          N/A         01:08:26
      Yes 3        xe17      No        20.1.1.1  -     -          -
L>     4.4.4.4/32    2         25        -          -     -          -
      -          N/A         01:08:26
      Yes 3        cd0/2     No        30.1.1.2  -     -          -
L>     5.5.5.5/32    3         28        -          -     -          -
      -          N/A         01:08:26
      Yes 24323    cd0/2     No        30.1.1.2  -     -          -
L>     6.6.6.6/32    4         36        -          -     -          -
      -          N/A         01:08:26
      Yes 24324    xe17      No        20.1.1.1  -     -          -
      Yes 24324    cd0/2     No        30.1.1.2  -     -          -
L>     10.1.1.0/24   8         37        -          -     -          -
      -          N/A         01:08:26

```

```

      Yes    3          xe17      1      No      -      20.1.1.1      -      -
L> 40.1.1.0/24      5      25      -      -      -      -
      -          N/A      01:08:26      10      -
      Yes    3          cd0/2      No      -      30.1.1.2      -      -
L> 50.1.1.0/24      6      33      -      -      -      -
      -          N/A      01:08:26      32      -
      Yes    24327      cd0/2      No      -      30.1.1.2      -      -
L> 60.1.1.0/24      9      40      -      -      -      -
      -          N/A      01:08:26      39      -
      Yes    24328      xe17      No      -      20.1.1.1      -      -

```

PE2#show mpls ftn-table

Primary FTN entry with FEC: 1.1.1.1/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:08:28, UpTime: 01:08:28, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 22 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 22, owner: LDP, Stale: NO, refcount: 2, out intf: xe17, out label: 24320

Nexthop addr: 20.1.1.1 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 2.2.2.2/32, id: 7, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:08:28, UpTime: 01:08:28, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 5, out intf: xe17, out label: 3

Nexthop addr: 20.1.1.1 cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 2, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:08:28, UpTime: 01:08:28, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 10 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 10, owner: N/A, Stale: NO, refcount: 5, out intf: cd0/2, out label: 3

Nexthop addr: 30.1.1.2 cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 3, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:08:28, UpTime: 01:08:28, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 27 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 27, owner: LDP, Stale: NO, refcount: 2, out intf: cd0/2, out label: 24323

Nexthop addr: 30.1.1.2 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 6.6.6.6/32, id: 4, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:08:28, UpTime: 01:08:28, LastUpdate: 01:08:28

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

```

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 35 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 35, owner: LDP, Stale: NO, refcount: 1, out intf: xe17, out label: 24324
  Nexthop addr: 20.1.1.1      cross connect ix: 8, op code: Push

  Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 30 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 30, owner: LDP, Stale: NO, refcount: 1, out intf: cd0/2, out label: 24324
  Nexthop addr: 30.1.1.2      cross connect ix: 8, op code: Push

Primary FTN entry with FEC: 10.1.1.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:08:28, UpTime: 01:08:28, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 5, out intf: xe17, out label: 3
  Nexthop addr: 20.1.1.1      cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 40.1.1.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:08:28, UpTime: 01:08:28, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 10 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 10, owner: N/A, Stale: NO, refcount: 5, out intf: cd0/2, out label: 3
  Nexthop addr: 30.1.1.2      cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 50.1.1.0/24, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:08:28, UpTime: 01:08:28, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 32 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 32, owner: LDP, Stale: NO, refcount: 2, out intf: cd0/2, out label: 24327
  Nexthop addr: 30.1.1.2      cross connect ix: 7, op code: Push

Primary FTN entry with FEC: 60.1.1.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:08:28, UpTime: 01:08:28, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 39 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 39, owner: LDP, Stale: NO, refcount: 2, out intf: xe17, out label: 24328
  Nexthop addr: 20.1.1.1      cross connect ix: 9, op code: Push

```

PE3

```

PE3#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

FTN-ECMP LDP: Disabled, SR: Disabled

| Code Intf | FEC ELC | Nexthop | FTN-ID | Nhlfe-ID Algo-Num | Tunnel-ID UpTime | Pri | Out-Label | Out- |
|--------------|-------------|---------|----------|----------------------|---------------------|-----|-----------|------|
| L> | 1.1.1.1/32 | 1 | 11 | - | - | - | - | - |
| | - | N/A | 01:09:10 | 3 | - | - | - | - |
| | Yes | 24320 | xe23 | No | 30.1.1.1 | - | - | - |
| | Yes | 24320 | xe2 | No | 40.1.1.2 | - | - | - |
| L> | 2.2.2.2/32 | 2 | 15 | - | - | - | - | - |
| | - | N/A | 01:09:10 | 7 | - | - | - | - |
| | Yes | 24321 | xe23 | No | 30.1.1.1 | - | - | - |
| L> | 3.3.3.3/32 | 3 | 28 | - | - | - | - | - |
| | - | N/A | 01:09:10 | 5 | - | - | - | - |
| | Yes | 3 | xe23 | No | 30.1.1.1 | - | - | - |
| L> | 5.5.5.5/32 | 6 | 35 | - | - | - | - | - |
| | - | N/A | 01:09:10 | 16 | - | - | - | - |
| | Yes | 3 | xe2 | No | 40.1.1.2 | - | - | - |
| L> | 6.6.6.6/32 | 7 | 38 | - | - | - | - | - |
| | - | N/A | 01:09:10 | 37 | - | - | - | - |
| | Yes | 24324 | xe2 | No | 40.1.1.2 | - | - | - |
| L> | 10.1.1.0/24 | 4 | 32 | - | - | - | - | - |
| | - | N/A | 01:09:10 | 31 | - | - | - | - |
| | Yes | 24325 | xe23 | No | 30.1.1.1 | - | - | - |
| L> | 20.1.1.0/24 | 5 | 28 | - | - | - | - | - |
| | - | N/A | 01:09:10 | 5 | - | - | - | - |
| | Yes | 3 | xe23 | No | 30.1.1.1 | - | - | - |
| L> | 50.1.1.0/24 | 8 | 35 | - | - | - | - | - |
| | - | N/A | 01:09:10 | 16 | - | - | - | - |
| | Yes | 3 | xe2 | No | 40.1.1.2 | - | - | - |
| L> | 60.1.1.0/24 | 9 | 40 | - | - | - | - | - |
| | - | N/A | 01:09:10 | 20 | - | - | - | - |
| | Yes | 24328 | xe2 | No | 40.1.1.2 | - | - | - |

PE3#show mpls ftn-table

Primary FTN entry with FEC: 1.1.1.1/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:12, UpTime: 01:09:12, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number: 0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 3 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 3, owner: LDP, Stale: NO, refcount: 3, out intf: xe23, out label: 24320
Nexthop addr: 30.1.1.1 cross connect ix: 2, op code: Push

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 34 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 34, owner: LDP, Stale: NO, refcount: 2, out intf: xe2, out label: 24320
Nexthop addr: 40.1.1.2 cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 2.2.2.2/32, id: 2, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:12, UpTime: 01:09:12, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number: 0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 3, out intf: xe23, out label: 24321

NextHop addr: 30.1.1.1 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 3, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:12, UpTime: 01:09:12, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 5 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 5, owner: N/A, Stale: NO, refcount: 7, out intf: xe23, out label: 3

NextHop addr: 30.1.1.1 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 6, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:12, UpTime: 01:09:12, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 16 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 7, out intf: xe2, out label: 3

NextHop addr: 40.1.1.2 cross connect ix: 7, op code: Push

Primary FTN entry with FEC: 6.6.6.6/32, id: 7, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:12, UpTime: 01:09:12, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 37 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 37, owner: LDP, Stale: NO, refcount: 2, out intf: xe2, out label: 24324

NextHop addr: 40.1.1.2 cross connect ix: 9, op code: Push

Primary FTN entry with FEC: 10.1.1.0/24, id: 4, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:12, UpTime: 01:09:12, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 31 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 31, owner: LDP, Stale: NO, refcount: 2, out intf: xe23, out label: 24325

NextHop addr: 30.1.1.1 cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 20.1.1.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:12, UpTime: 01:09:12, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 5 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 5, owner: N/A, Stale: NO, refcount: 7, out intf: xe23, out label: 3

NextHop addr: 30.1.1.1 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 50.1.1.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:12, UpTime: 01:09:12, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 16 refcount: 1
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 7, out intf: xe2, out label: 3
 Nexthop addr: 40.1.1.2 cross connect ix: 7, op code: Push

Primary FTN entry with FEC: 60.1.1.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:12, UpTime: 01:09:12, LastUpdate: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number: 0
 Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
 Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 20 refcount: 1
 Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 20, owner: LDP, Stale: NO, refcount: 3, out intf: xe2, out label: 24328
 Nexthop addr: 40.1.1.2 cross connect ix: 6, op code: Push

P2

P2#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|-------------|---------|----------|-----------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 6 | 32 | - | - | - | - |
| | - | N/A | 01:09:30 | | | | |
| | Yes | 24320 | ge6 | No | | 50.1.1.2 | - |
| L> | 2.2.2.2/32 | 1 | 3 | - | - | - | - |
| | - | N/A | 01:09:31 | | | | |
| | Yes | 24321 | ge4 | No | | 40.1.1.1 | - |
| | Yes | 24321 | ge6 | No | | 50.1.1.2 | - |
| L> | 3.3.3.3/32 | 2 | 9 | - | - | - | - |
| | - | N/A | 01:09:31 | | | | |
| | Yes | 24322 | ge4 | No | | 40.1.1.1 | - |
| L> | 4.4.4.4/32 | 4 | 29 | - | - | - | - |
| | - | N/A | 01:09:31 | | | | |
| | Yes | 3 | ge4 | No | | 40.1.1.1 | - |
| L> | 6.6.6.6/32 | 7 | 35 | - | - | - | - |
| | - | N/A | 01:09:30 | | | | |
| | Yes | 3 | ge6 | No | | 50.1.1.2 | - |
| L> | 10.1.1.0/24 | 8 | 38 | - | - | - | - |
| | - | N/A | 01:09:30 | | | | |
| | Yes | 24325 | ge6 | No | | 50.1.1.2 | - |
| L> | 20.1.1.0/24 | 3 | 14 | - | - | - | - |
| | - | N/A | 01:09:31 | | | | |
| | Yes | 24326 | ge4 | No | | 40.1.1.1 | - |
| L> | 30.1.1.0/24 | 5 | 29 | - | - | - | - |
| | - | N/A | 01:09:31 | | | | |
| | Yes | 3 | ge4 | No | | 40.1.1.1 | - |
| L> | 60.1.1.0/24 | 9 | 35 | - | - | - | - |
| | - | N/A | 01:09:30 | | | | |
| | Yes | 3 | ge6 | No | | 50.1.1.2 | - |

```
P2#show mpls ftn-table
Primary FTN entry with FEC: 1.1.1.1/32, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:32, UpTime: 01:09:32, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 31 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 31, owner: LDP, Stale: NO, refcount: 2, out intf: ge6, out label: 24320
      Nexthop addr: 50.1.1.2      cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:33, UpTime: 01:09:33, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 5 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 5, owner: LDP, Stale: NO, refcount: 2, out intf: ge4, out label: 24321
      Nexthop addr: 40.1.1.1      cross connect ix: 3, op code: Push

    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 34 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 34, owner: LDP, Stale: NO, refcount: 1, out intf: ge6, out label: 24321
      Nexthop addr: 50.1.1.2      cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:33, UpTime: 01:09:33, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1
      Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 4, owner: LDP, Stale: NO, refcount: 2, out intf: ge4, out label: 24322
      Nexthop addr: 40.1.1.1      cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:33, UpTime: 01:09:33, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 5, out intf: ge4, out label: 3
      Nexthop addr: 40.1.1.1      cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 6.6.6.6/32, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:32, UpTime: 01:09:32, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 25 refcount: 1
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      Out-segment with ix: 25, owner: N/A, Stale: NO, refcount: 5, out intf: ge6, out label: 3
      Nexthop addr: 50.1.1.2      cross connect ix: 8, op code: Push

Primary FTN entry with FEC: 10.1.1.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:32, UpTime: 01:09:32, LastUpdate: N/A
```

```

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 37 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 37, owner: LDP, Stale: NO, refcount: 2, out intf: ge6, out label: 24325
Nexthop addr: 50.1.1.2 cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 20.1.1.0/24, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:09:33, UpTime: 01:09:33, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 13 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 13, owner: LDP, Stale: NO, refcount: 2, out intf: ge4, out label: 24326
Nexthop addr: 40.1.1.1 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 30.1.1.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:09:33, UpTime: 01:09:33, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 5, out intf: ge4, out label: 3
Nexthop addr: 40.1.1.1 cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 60.1.1.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 01:09:32, UpTime: 01:09:32, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 25 refcount: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 25, owner: N/A, Stale: NO, refcount: 5, out intf: ge6, out label: 3
Nexthop addr: 50.1.1.2 cross connect ix: 8, op code: Push

```

PE4

```

PE4#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code   FEC           FTN-ID   Nhlfe-ID  Tunnel-ID   Pri   Out-Label   Out-
Intf   ELC           Nexthop   Algo-Num   UpTime
L> 1.1.1.1/32    6        32        -          -      -          -
    -           N/A      01:09:55
    Yes 3        xe2       No        60.1.1.2   -      -          -
L> 2.2.2.2/32    7        35        -          -      -          -
    -           N/A      01:09:54
    Yes 24320     xe2       No        60.1.1.2   -      -          -
L> 3.3.3.3/32    1        38        -          -      -          -
    -           N/A      01:09:55
    Yes 24321     xe2       No        60.1.1.2   -      -          -

```



```

                22      -      50.1.1.1      -      -
L>      Yes      24322      xe6      No      -      -      -
      4.4.4.4/32      2      25      -      -      -
      -      N/A      01:09:55      24      -
                24      -      50.1.1.1      -      -
L>      Yes      24323      xe6      No      -      -      -
      5.5.5.5/32      3      27      -      -      -
      -      N/A      01:09:55      16      -
                16      -      50.1.1.1      -      -
L>      Yes      3      xe6      No      -      -      -
      10.1.1.0/24      8      32      -      -      -
      -      N/A      01:09:54      1      -
                1      -      60.1.1.2      -      -
L>      Yes      3      xe2      No      -      -      -
      20.1.1.0/24      9      40      -      -      -
      -      N/A      01:09:54      39      -
                39      -      60.1.1.2      -      -
L>      Yes      24325      xe2      No      -      -      -
      30.1.1.0/24      4      30      -      -      -
      -      N/A      01:09:55      29      -
                29      -      50.1.1.1      -      -
L>      Yes      24327      xe6      No      -      -      -
      40.1.1.0/24      5      27      -      -      -
      -      N/A      01:09:55      16      -
                16      -      50.1.1.1      -      -
      Yes      3      xe6      No      -      -

```

PE4#show mpls ftn-table

Primary FTN entry with FEC: 1.1.1.1/32, id: 6, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:57, UpTime: 01:09:57, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 5, out intf: xe2, out label: 3

Nexthop addr: 60.1.1.2 cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 2.2.2.2/32, id: 7, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:56, UpTime: 01:09:56, LastUpdate: N/A

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 34 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 34, owner: LDP, Stale: NO, refcount: 2, out intf: xe2, out label: 24320

Nexthop addr: 60.1.1.2 cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 01:09:57, UpTime: 01:09:57, LastUpdate: 01:09:56

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 37 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 37, owner: LDP, Stale: NO, refcount: 1, out intf: xe2, out label: 24321

Nexthop addr: 60.1.1.2 cross connect ix: 8, op code: Push

Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 22 refcount: 1

Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 22, owner: LDP, Stale: NO, refcount: 1, out intf: xe6, out label: 24322

Nexthop addr: 50.1.1.1 cross connect ix: 8, op code: Push

```
Primary FTN entry with FEC: 4.4.4.4/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:57, UpTime: 01:09:57, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 24 refcount: 1
    Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 24, owner: LDP, Stale: NO, refcount: 2, out intf: xe6, out label: 24323
    Nexthop addr: 50.1.1.1      cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:57, UpTime: 01:09:57, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 16 refcount: 1
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 5, out intf: xe6, out label: 3
    Nexthop addr: 50.1.1.1      cross connect ix: 7, op code: Push

Primary FTN entry with FEC: 10.1.1.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:56, UpTime: 01:09:56, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1 refcount: 1
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 5, out intf: xe2, out label: 3
    Nexthop addr: 60.1.1.2      cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 20.1.1.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:56, UpTime: 01:09:56, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 39 refcount: 1
    Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 39, owner: LDP, Stale: NO, refcount: 2, out intf: xe2, out label: 24325
    Nexthop addr: 60.1.1.2      cross connect ix: 9, op code: Push

Primary FTN entry with FEC: 30.1.1.0/24, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:57, UpTime: 01:09:57, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 29 refcount: 1
    Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 29, owner: LDP, Stale: NO, refcount: 2, out intf: xe6, out label: 24327
    Nexthop addr: 50.1.1.1      cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 40.1.1.0/24, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 01:09:57, UpTime: 01:09:57, LastUpdate: N/A
  Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:0
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 16 refcount: 1
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

Out-segment with ix: 16, owner: N/A, Stale: NO, refcount: 5, out intf: xe6, out label: 3
Nexthop addr: 50.1.1.1 cross connect ix: 7, op code: Push

LDP MD5 Password For Auto Targeted Sessions

RFC 5286 describes a basic mechanism using Loop-Free Alternates (LFAs) that provides good repair coverage in many topologies, especially those that are highly meshed.

However, some topologies, notably ring-based topologies, are not well protected by LFAs alone. This is because there is no neighbor of the Point of Local Repair (PLR) that has a cost to the destination via a path that does not traverse the failure that is cheaper than the cost to the destination via the failure.

When LDP RLFA creates a virtual path in the network to provide an alternate path, it uses MPLS labels distributed by a targeted session between the local node and the PLR (PQ node). The session is established without any MD5 Password protection mechanism, as it is only available for non-automatic targeted sessions. A set of configurable options will be provided to associate MD5 passwords with the auto-targeted sessions.

Configuration


The following configuration example demonstrates how to enable LDP authentication for an auto-targeted peer session between PE1 and PE2 using their loopback interfaces.

Topology

The topology illustrates **provider edge (PE) routers**, PE1, PE2, PE3, and PE4, which are connected through a **core network (P1 and P2)** consisting of intermediate P and PE routers. All PE routers are configured to establish an **LDP auto-targeted session**. The diagram highlights the **use of LDP over targeted sessions** between the loopback interfaces of the PE routers. It also demonstrates how **LDP authentication** can be applied to this auto-discovered targeted session, ensuring secure label distribution across the network core without manually defining LDP peers.

The following topology shows the configuration required to enable the RLFA feature.

Figure 15. RLFA Topology



PE1

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable the LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To allow the creation of a TLDP session dynamically |
| (config-router)#neighbor auto-targeted auth md5 password plain-text test2 | Configure MD5 authentication for auto-targeted peers |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit the LDP process |
| (config)#router rsvp | Enable RSVP |
| (config-router)#srlg-disjoint forced | Configure srlg to enable MPLS-TE |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)#interface ce6/2 | Enter interface mode. |
| (config-if)#ip address 10.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure the OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the ce6/2 interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable Mpls-TE |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe29 | Enter interface mode. |
| (config-if)#ip address 60.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable mpls-TE |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the xe29 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure the IP address of the interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Create an OSPF instance |
| (config-router)# ospf router-id 1.1.1.1 | Configure router ID |

| | |
|---|--|
| (config-router)# bfd all-interfaces | Configure bfd |
| (config-router)# timers spf exp 50 50 | Configure the OSPF timers |
| (config-router)# timers throttle lsa all 0 1 1 | Configure the OSPF timer LSA throttle |
| (config-router)# timers lsa arrival 1 | Configure the OSPF timer for LSA arrival |
| (config-router)# fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp | Configure Remote LFA to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| (config-router)# network 1.1.1.1/32 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 10.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 60.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#end | Exit router mode. |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |

P1

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable the LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To allow the creation of a TLDP session dynamically |
| (config-router)#neighbor auto-targeted auth md5 password plain-text test2 | Configure MD5 authentication for auto-targeted peers |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit the LDP process |
| (config)#router rsvp | Enable RSVP |
| (config-router)#srlg-disjoint forced | Configure srlg to enable MPLS-TE |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)#interface xe8 | Enter interface mode. |
| (config-if)#ip address 10.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure the OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the xe8 interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable Mpls-TE |
| (config-if)#commit | Commit the candidate configuration to the running |

| | |
|---|---|
| | configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe17 | Enter interface mode. |
| (config-if)#ip address 20.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable mpls-TE |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the xe17 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure the IP address of the interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Create an OSPF instance |
| (config-router)# ospf router-id 2.2.2.2 | Configure router ID |
| (config-router)# bfd all-interfaces | Configure bfd |
| (config-router)# timers spf exp 50 50 | Configure the OSPF timers |
| (config-router)# timers throttle lsa all 0 1 1 | Configure the OSPF timer LSA throttle |
| (config-router)# timers lsa arrival 1 | Configure the OSPF timer for LSA arrival |
| (config-router)# network 2.2.2.2/32 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 10.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 20.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | Exit router mode. |

PE2

| | |
|---------------------------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable the LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To allow the creation a TLDP session dynamically |

| | |
|---|---|
| (config-router)#neighbor auto-targeted auth md5 password plain-text test2 | Configure MD5 authentication for auto-targeted peers |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit the LDP process |
| (config)#router rsvp | Enable RSPVP |
| (config-router)#srlg-disjoint forced | Configure srlg to enable MPLS-TE |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)#interface xe17 | Enter interface mode. |
| (config-if)#ip address 20.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure the OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the xe17 interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable Mpls-TE |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface cd0/2 | Enter interface mode. |
| (config-if)#ip address 30.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable mpls-TE |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the cd0/2 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure the IP address of the interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Create an OSPF instance |
| (config-router)# ospf router-id 3.3.3.3 | Configure router ID |
| (config-router)# bfd all-interfaces | Configure bfd |
| (config-router)# timers spf exp 50 50 | Configure the OSPF timers |
| (config-router)# timers throttle lsa all 0 1 1 | Configure the OSPF timer LSA throttle |
| (config-router)# timers lsa arrival 1 | Configure the OSPF timer for LSA arrival |
| (config-router)# network 3.3.3.3/32 area 0.0.0.0 | Configure the network command to advertise the prefixes |

| | |
|---|---|
| (config-router)# network 20.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 30.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | Exit router mode. |

PE3

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable the LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To allow the creation of a TLDP session dynamically |
| (config-router)#neighbor auto-targeted auth md5 password plain-text test2 | Configure MD5 authentication for auto-targeted peers |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit the LDP process |
| (config)#router rsvp | Enable RSVP |
| (config-router)#srlg-disjoint forced | Configure srlg to enable MPLS-TE |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)#interface xe23 | Enter interface mode. |
| (config-if)#ip address 30.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure the OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the xe23 interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable Mpls-TE |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 40.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |

| | |
|---|---|
| (config-if)#mpls traffic-eng srlg 11 | Enable mpls-TE |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the xe2 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure the IP address of the interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Create an OSPF instance |
| (config-router)# ospf router-id 4.4.4.4 | Configure router ID |
| (config-router)# bfd all-interfaces | Configure bfd |
| (config-router)# timers spf exp 50 50 | Configure the OSPF timers |
| (config-router)# timers throttle lsa all 0 1 1 | Configure the OSPF timer LSA throttle |
| (config-router)# timers lsa arrival 1 | Configure the OSPF timer for LSA arrival |
| (config-router)# network 4.4.4.4/32 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 40.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 30.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | Exit router mode. |

P2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable the LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To allow the creation of a TLDP session dynamically |
| (config-router)#neighbor auto-targeted auth md5 password plain-text test2 | Configure MD5 authentication for auto-targeted peers |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit the LDP process |
| (config)#router rsvp | Enable RSVP |

| | |
|---|---|
| (config-router)#srlg-disjoint forced | Configure srlg to enable MPLS-TE |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)#interface ge4 | Enter interface mode. |
| (config-if)#ip address 40.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure the OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the ge4 interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable Mpls-TE |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface ge6 | Enter interface mode. |
| (config-if)#ip address 50.1.1.1/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable mpls-TE |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the ge6 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 5.5.5.5/32 secondary | Configure the IP address of the interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Create an OSPF instance |
| (config-router)# ospf router-id 5.5.5.5 | Configure router ID |
| (config-router)# bfd all-interfaces | Configure bfd |
| (config-router)# timers spf exp 50 50 | Configure the OSPF timers |
| (config-router)# timers throttle lsa all 0 1 1 | Configure the OSPF timer LSA throttle |
| (config-router)# timers lsa arrival 1 | Configure the OSPF timer for LSA arrival |
| (config-router)# network 5.5.5.5/32 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 50.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 40.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure bfd interval globally |

| | |
|-----------------|---|
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | Exit router mode. |

PE4

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enable the LDP process |
| (config-router)#fast-reroute | Enable LDP FRR |
| (config-router)#auto-targeted-session | To allow the creation of a TLDP session dynamically |
| (config-router)#neighbor auto-targeted auth md5 password plain-text test2 | Configure MD5 authentication for auto-targeted peers |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config-router)#exit | Exit the LDP process |
| (config)#router rsvp | Enable RSVP |
| (config-router)#srlg-disjoint forced | Configure srlg to enable MPLS-TE |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)#interface xe6 | Enter interface mode. |
| (config-if)#ip address 50.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure the OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the xe6 interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable Mpls-TE |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 60.1.1.2/24 | Configure the IP address of the interface. |
| (config-if)#ip ospf cost 10 | Configure OSPF cost |
| (config-if)#label-switching | Enable label-switching on the interface |
| (config-if)#mpls traffic-eng srlg 11 | Enable mpls-TE |
| (config-if)#enable-ldp ipv4 | Enable ldp process on the xe2 interface |
| (config-if)#commit | Commit the candidate configuration to the running configuration |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 6.6.6.6/32 secondary | Configure the IP address of the interface |

| | |
|---|---|
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Create an OSPF instance |
| (config-router)# ospf router-id 6.6.6.6 | Configure router ID |
| (config-router)# bfd all-interfaces | Configure BFD |
| (config-router)# timers spf exp 50 50 | Configure the OSPF timers |
| (config-router)# timers throttle lsa all 0 1 1 | Configure the OSPF timer LSA throttle |
| (config-router)# timers lsa arrival 1 | Configure the OSPF timer for LSA arrival |
| (config-router)# network 6.6.6.6/32 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 50.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)# network 60.1.1.0/24 area 0.0.0.0 | Configure the network command to advertise the prefixes |
| (config-router)#commit | Commit the candidate configuration to the running configuration |
| (config)# bfd interval 3 minrx 3 multiplier 3 | Configure the BFD interval globally |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | Exit router mode. |

Validation

PE1

Check LDP neighborship before enabling RLFA

```
PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State        KeepAlive  UpTime
-----
2.2.2.2      ce6/2      Passive    OPERATIONAL 30          00:29:22
6.6.6.6      xe29       Passive    OPERATIONAL 30          00:05:30
```

Check the output of "show ip ospf neighbors" to verify that ospf adjacency is up.

```
PE1#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID    Pri  State           Dead Time   Address        Interface      Instance ID
-----
2.2.2.2        1    Full/Backup     00:00:36   10.1.1.2       ce6/2          0
6.6.6.6        1    Full/Backup     00:00:33   60.1.1.2       xe29           0
```

Check the ospf route installation in the ospf table and RIB table.
PE1#show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP

```

O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN,
v - vrf leaked
* - candidate default

IP Route Table for VRF "default"
C      1.1.1.1/32 is directly connected, lo, installed 00:36:44, last update 00:36:44 ago
O      2.2.2.2/32 [110/11] via 10.1.1.2, ce6/2, installed 00:30:55, last update 00:30:55 ago
O      3.3.3.3/32 [110/21] via 10.1.1.2, ce6/2, installed 00:26:57, last update 00:26:57 ago
O      4.4.4.4/32 [110/31] via 60.1.1.2, xe29, installed 00:23:06, last update 00:06:21 ago
       [110/31] via 10.1.1.2, ce6/2
O      5.5.5.5/32 [110/21] via 60.1.1.2, xe29, installed 00:06:21, last update 00:06:21 ago
O      6.6.6.6/32 [110/11] via 60.1.1.2, xe29, installed 00:06:26, last update 00:06:26 ago
C      10.1.1.0/24 is directly connected, ce6/2, installed 00:39:08, last update 00:39:08 ago
O      20.1.1.0/24 [110/20] via 10.1.1.2, ce6/2, installed 00:30:55, last update 00:30:55 ago
O      30.1.1.0/24 [110/30] via 10.1.1.2, ce6/2, installed 00:26:57, last update 00:26:57 ago
O      40.1.1.0/24 [110/30] via 60.1.1.2, xe29, installed 00:06:21, last update 00:06:21 ago
O      50.1.1.0/24 [110/20] via 60.1.1.2, xe29, installed 00:06:26, last update 00:06:26 ago
C      60.1.1.0/24 is directly connected, xe29, installed 00:36:57, last update 00:36:57 ago
C      127.0.0.0/8 is directly connected, lo, installed 02:35:26, last update 02:35:26 ago

```

Gateway of last resort is not set

```
PE1#show ip ospf route
```

```

OSPF process 1:
Codes: C - connected, D - Discard, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
OSPF LFA attributes:
P - Primary, SP - Secondary-Path, LP - Link Protecting,
NP - Node Protecting, BID - Broadcast Link Protecting
DP - Downstream Protecting

C 1.1.1.1/32 [1] is directly connected, lo, Area 0.0.0.0
O 2.2.2.2/32 [11] via 10.1.1.2, ce6/2, Area 0.0.0.0
O 3.3.3.3/32 [21] via 10.1.1.2, ce6/2, Area 0.0.0.0
O 4.4.4.4/32 [31] via 10.1.1.2, ce6/2, Area 0.0.0.0
  via 60.1.1.2, xe29, Area 0.0.0.0
O 5.5.5.5/32 [21] via 60.1.1.2, xe29, Area 0.0.0.0
O 6.6.6.6/32 [11] via 60.1.1.2, xe29, Area 0.0.0.0
C 10.1.1.0/24 [10] is directly connected, ce6/2, Area 0.0.0.0
O 20.1.1.0/24 [20] via 10.1.1.2, ce6/2, Area 0.0.0.0
O 30.1.1.0/24 [30] via 10.1.1.2, ce6/2, Area 0.0.0.0
O 40.1.1.0/24 [30] via 60.1.1.2, xe29, Area 0.0.0.0
O 50.1.1.0/24 [20] via 60.1.1.2, xe29, Area 0.0.0.0
C 60.1.1.0/24 [10] is directly connected, xe29, Area 0.0.0.0

```

Verify parameters related to LDP using “show ldp” command

```

PE1#show ldp
Router ID           : 10.1.1.1
LDP Version         : 1
Fast-reroute Per-prefix : Enabled
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode  : Liberal
Label Control Mode    : Independent
Instance Loop Detection : Off
Request Retry         : Off
Propagate Release      : Disabled
Graceful Restart       : Disabled
Hello Interval        : 5

```

```

Targeted Hello Interval : 15
Hold time                : 15
Targeted Hold time      : 45
Keepalive Interval      : 10
Keepalive Timeout       : 30
Request retry Timeout   : 5
Auto Targeted Hello Receipt : Enabled
Transport Address data  :
  Labelspace 0          : 1.1.1.1 (in use)
Import BGP routes       : No
Entropy-Label Capability : Disabled
Prefer-Tunnel-In-Tunnel : Disabled

```

Verify ospf LFA and RLFA backup computed paths for Primary Paths

```

PE1#show ip ospf route fast-reroute

OSPF process 1:
Codes: C - connected, D - Discard, O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
OSPF LFA attributes:
      P - Primary, SP - Secondary-Path, LP - Link Protecting,
      NP - Node Protecting, BID - Broadcast Link Protecting
      DP - Downstream Protecting

O 2.2.2.2/32 [11] via 10.1.1.2, ce6/2, Area 0.0.0.0
  Remote FRR path : via 4.4.4.4, via 60.1.1.2, xe29, Area 0.0.0.0
  Attributes      :
  Metric          : [51] ,LP
O 3.3.3.3/32 [21] via 10.1.1.2, ce6/2, Area 0.0.0.0
  Remote FRR path : via 4.4.4.4, via 60.1.1.2, xe29, Area 0.0.0.0
  Attributes      :
  Metric          : [41] ,LP ,NP ,DP
O 4.4.4.4/32 [31] via 10.1.1.2, ce6/2, Area 0.0.0.0
  Backup path     : via 60.1.1.2, xe29, Area 0.0.0.0
  Attributes      :
  Metric          : [31] ,P ,NP ,BID ,DP
via 60.1.1.2, xe29, Area 0.0.0.0
  Backup path     : via 10.1.1.2, ce6/2, Area 0.0.0.0
  Attributes      :
  Metric          : [31] ,P ,NP ,BID ,DP
O 5.5.5.5/32 [21] via 60.1.1.2, xe29, Area 0.0.0.0
  Remote FRR path : via 4.4.4.4, via 10.1.1.2, ce6/2, Area 0.0.0.0
  Attributes      :
  Metric          : [41] ,LP ,NP ,DP
O 6.6.6.6/32 [11] via 60.1.1.2, xe29, Area 0.0.0.0
  Remote FRR path : via 4.4.4.4, via 10.1.1.2, ce6/2, Area 0.0.0.0
  Attributes      :
  Metric          : [51] ,LP
O 20.1.1.0/24 [20] via 10.1.1.2, ce6/2, Area 0.0.0.0
  Remote FRR path : via 4.4.4.4, via 60.1.1.2, xe29, Area 0.0.0.0
  Attributes      :
  Metric          : [50] ,LP ,NP
O 30.1.1.0/24 [30] via 10.1.1.2, ce6/2, Area 0.0.0.0
  Backup path     : via 60.1.1.2, xe29, Area 0.0.0.0
  Attributes      :
  Metric          : [40] ,SP ,NP ,BID
O 40.1.1.0/24 [30] via 60.1.1.2, xe29, Area 0.0.0.0
  Backup path     : via 10.1.1.2, ce6/2, Area 0.0.0.0
  Attributes      :
  Metric          : [40] ,SP ,NP ,BID
O 50.1.1.0/24 [20] via 60.1.1.2, xe29, Area 0.0.0.0
  Remote FRR path : via 4.4.4.4, via 10.1.1.2, ce6/2, Area 0.0.0.0
  Attributes      :
  Metric          : [50] ,LP ,NP

```

Verify PQ node which is near to source is selected and Target-LDP session is established with PQ node using below commands

```
PE1#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | ce6/2 | Passive | OPERATIONAL | 30 | 00:31:05 |
| | 6.6.6.6 | xe29 | Passive | OPERATIONAL | 30 | 00:07:13 |
| | 4.4.4.4 | ce6/2 | Passive | OPERATIONAL | 30 | 00:07:14 |

To verify which password is enabled

```
PE1#show ldp session 4.4.4.4
```

```
Session state           : OPERATIONAL
Session role            : Passive
TCP Connection          : Established
IP Address for TCP      : 4.4.4.4
Interface being used    : ce6/2
Peer LDP ID             : 30.1.1.2:0
Last Down Reason        : None
Preferred Peer LDP Password : test2
Current Peer LDP Password : test2
Authentication type     : MD5
Adjacencies             : 4.4.4.4
Advertisement mode       : Downstream Unsolicited
Label retention mode     : Liberal
Graceful Restart        : Not Capable
Keepalive Timeout       : 30
Reconnect Interval      : 15
Configured TCP MSS      : Not configured
Applied TCP MSS         : 1440
Preferred TCP MSS        : NA
Address List received    : 4.4.4.4
                        30.1.1.2
                        40.1.1.1
```

| Received Labels : | Fec | Label | Maps To |
|-------------------|------------------|-------|---------|
| | IPV4:60.1.1.0/24 | 24965 | none |
| | IPV4:40.1.1.0/24 | 24971 | none |
| | IPV4:30.1.1.0/24 | 24970 | none |
| | IPV4:10.1.1.0/24 | 24963 | none |
| | IPV4:4.4.4.4/32 | 24969 | none |
| | IPV4:1.1.1.1/32 | 24960 | none |
| | IPV4:50.1.1.0/24 | 24967 | 24967 |
| | IPV4:50.1.1.0/24 | 24967 | 24967 |
| | IPV4:50.1.1.0/24 | 24967 | 24967 |
| | IPV4:20.1.1.0/24 | 24964 | 24961 |
| | IPV4:20.1.1.0/24 | 24964 | 24961 |
| | IPV4:20.1.1.0/24 | 24964 | 24961 |
| | IPV4:6.6.6.6/32 | 24968 | 24968 |
| | IPV4:6.6.6.6/32 | 24968 | 24968 |
| | IPV4:6.6.6.6/32 | 24968 | 24968 |
| | IPV4:5.5.5.5/32 | 24966 | 24966 |
| | IPV4:5.5.5.5/32 | 24966 | 24966 |
| | IPV4:5.5.5.5/32 | 24966 | 24966 |
| | IPV4:3.3.3.3/32 | 24962 | 24962 |
| | IPV4:3.3.3.3/32 | 24962 | 24962 |
| | IPV4:3.3.3.3/32 | 24962 | 24962 |
| | IPV4:2.2.2.2/32 | 24961 | 24960 |
| | IPV4:2.2.2.2/32 | 24961 | 24960 |
| | IPV4:2.2.2.2/32 | 24961 | 24960 |

| Sent Labels : | Fec | Label | Maps To |
|---------------|-----|-------|---------|
|---------------|-----|-------|---------|

| | | |
|------------------|-------|-----------|
| IPv4:60.1.1.0/24 | 24971 | none |
| IPv4:50.1.1.0/24 | 24967 | 24967 |
| IPv4:50.1.1.0/24 | 24967 | impl-null |
| IPv4:40.1.1.0/24 | 24965 | 24965 |
| IPv4:40.1.1.0/24 | 24965 | 24968 |
| IPv4:30.1.1.0/24 | 24963 | 24963 |
| IPv4:30.1.1.0/24 | 24963 | 24967 |
| IPv4:20.1.1.0/24 | 24961 | impl-null |
| IPv4:20.1.1.0/24 | 24961 | 24964 |
| IPv4:10.1.1.0/24 | 24970 | none |
| IPv4:6.6.6.6/32 | 24968 | 24968 |
| IPv4:6.6.6.6/32 | 24968 | impl-null |
| IPv4:5.5.5.5/32 | 24966 | 24966 |
| IPv4:5.5.5.5/32 | 24966 | 24964 |
| IPv4:4.4.4.4/32 | 24964 | 24964 |
| IPv4:4.4.4.4/32 | 24964 | 24963 |
| IPv4:3.3.3.3/32 | 24962 | 24962 |
| IPv4:3.3.3.3/32 | 24962 | 24962 |
| IPv4:2.2.2.2/32 | 24960 | impl-null |
| IPv4:2.2.2.2/32 | 24960 | 24961 |
| IPv4:1.1.1.1/32 | 24969 | none |

Verify that Primary and Backup FTN's are installed with labels in LDP RLFA route table

```
PE1#show ldp rlfa-routes
Codes: p - stale rLFA route
Fec          Primary-NH      Backup-NH      rLFA-Addr      Out-Intf  Outer-label  Inner-label  Owner
2.2.2.2      10.1.1.2      60.1.1.2      4.4.4.4        xe29      24963       24961      osp
f
3.3.3.3      10.1.1.2      60.1.1.2      4.4.4.4        xe29      24963       24962      osp
f
5.5.5.5      60.1.1.2      10.1.1.2      4.4.4.4        ce6/2     24964       24966      osp
f
6.6.6.6      60.1.1.2      10.1.1.2      4.4.4.4        ce6/2     24964       24968      osp
f
20.1.1.0     10.1.1.2      60.1.1.2      4.4.4.4        xe29      24963       24964      osp
f
50.1.1.0     60.1.1.2      10.1.1.2      4.4.4.4        ce6/2     24964       24967      osp
f
```

Verify that backup XC's calculated for primary FTN's in MPLS forwarding table. Verify the same in FTN table.

```
PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code  FEC          Nexthop      FTN-ID  Nhlfe-ID  Tunnel-ID  Pri  Out-Label  Out-
Intf  ELC          Algo-Num  UpTime
L> 2.2.2.2/32    1          22        -         -         -    -         -
-      N/A        00:33:46  1         -
Yes 3          ce6/2     No        10.1.1.2  -    -         -
No 24961       xe29      No        4.4.4.4  -    -         -
60.1.1.2, label 24963)
L> 3.3.3.3/32    3          24        -         -         -    -         -
-      N/A        00:30:46  3         -
Yes 24962       ce6/2     No        10.1.1.2  -    -         -
No 24962       xe29      No        4.4.4.4  -    -         -
(via
```

```

60.1.1.2, label 24963)
L> 4.4.4.4/32      5      20      -      -      -      -      -
      -      N/A      00:26:55
                        7
      Yes  24964      ce6/2      No      10.1.1.2      -      -
                        19
      No   24963      xe29      No      60.1.1.2      -      -
                        19
      Yes  24963      xe29      No      60.1.1.2      -      -
                        7
      No   24964      ce6/2      No      10.1.1.2      -      -
L> 5.5.5.5/32      6      30      -      -      -      -      -
      -      N/A      00:09:57
                        28
      Yes  24964      xe29      No      60.1.1.2      -      -
                        29
      No   24966      ce6/2      No      4.4.4.4      -      -

10.1.1.2, label 24964)
L> 6.6.6.6/32      7      35      -      -      -      -      -
      -      N/A      00:09:57
                        33
      Yes  3          xe29      No      60.1.1.2      -      -
                        34
      No   24968      ce6/2      No      4.4.4.4      -      -

10.1.1.2, label 24964)
L> 20.1.1.0/24     2      39      -      -      -      -      -
      -      N/A      00:33:46
                        1
      Yes  3          ce6/2      No      10.1.1.2      -      -
                        38
      No   24964      xe29      No      4.4.4.4      -      -

60.1.1.2, label 24963)
L> 30.1.1.0/24     4      42      -      -      -      -      -
      -      N/A      00:30:46
                        5
      Yes  24963      ce6/2      No      10.1.1.2      -      -
                        41
      No   24967      xe29      No      60.1.1.2      -      -
L> 40.1.1.0/24     8      46      -      -      -      -      -
      -      N/A      00:09:57
                        44
      Yes  24968      xe29      No      60.1.1.2      -      -
                        45
      No   24965      ce6/2      No      10.1.1.2      -      -
L> 50.1.1.0/24     9      50      -      -      -      -      -
      -      N/A      00:09:57
                        33
      Yes  3          xe29      No      60.1.1.2      -      -
                        49
      No   24967      ce6/2      No      4.4.4.4      -      -

10.1.1.2, label 24964)

```

PE1#show mpls ftn-table

Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed

CreateTime: 00:35:22, UpTime: 00:35:22, LastUpdate: 00:11:33

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none, Algorithm Number:0

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 4, out intf: ce6/2, out label: 3
 Nexthop addr: 10.1.1.2 cross connect ix: 1, op code: Push

```
Backup Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 21 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 21, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24961
Nexthop addr: 4.4.4.4 cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 00:32:22, UpTime: 00:32:22, LastUpdate: 00:11:33
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 3 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: LDP, Stale: NO, refcount: 2, out intf: ce6/2, out label: 24962
Nexthop addr: 10.1.1.2 cross connect ix: 2, op code: Push

Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 23 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 23, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24962
Nexthop addr: 4.4.4.4 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 4.4.4.4/32, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 00:28:31, UpTime: 00:28:31, LastUpdate: 00:11:33
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 7 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 4, out intf: ce6/2, out label: 24964
Nexthop addr: 10.1.1.2 cross connect ix: 4, op code: Push

Backup Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 19 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 19, owner: LDP, Stale: NO, refcount: 4, out intf: xe29, out label: 24963
Nexthop addr: 60.1.1.2 cross connect ix: 4, op code: Push

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 19 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 19, owner: LDP, Stale: NO, refcount: 4, out intf: xe29, out label: 24963
Nexthop addr: 60.1.1.2 cross connect ix: 4, op code: Push

Backup Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 7 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 7, owner: LDP, Stale: NO, refcount: 4, out intf: ce6/2, out label: 24964
Nexthop addr: 10.1.1.2 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 5.5.5.5/32, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 00:11:33, UpTime: 00:11:33, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 28 refcount: 1
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 28, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24964
Nexthop addr: 60.1.1.2 cross connect ix: 5, op code: Push

Backup Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 29 bypass ftn-ix: 0
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 29, owner: LDP, Stale: NO, refcount: 1, out intf: ce6/2, out label: 24966
Nexthop addr: 4.4.4.4 cross connect ix: 9, op code: Push

Primary FTN entry with FEC: 6.6.6.6/32, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
CreateTime: 00:11:33, UpTime: 00:11:33, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 33 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 33, owner: N/A, Stale: NO, refcount: 4, out intf: xe29, out label: 3
  Nexthop addr: 60.1.1.2      cross connect ix: 6, op code: Push

  Backup Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 34 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 34, owner: LDP, Stale: NO, refcount: 1, out intf: ce6/2, out label: 24968
  Nexthop addr: 4.4.4.4      cross connect ix: 10, op code: Push

Primary FTN entry with FEC: 20.1.1.0/24, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 00:35:22, UpTime: 00:35:22, LastUpdate: 00:11:33
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 4, out intf: ce6/2, out label: 3
  Nexthop addr: 10.1.1.2      cross connect ix: 1, op code: Push

  Backup Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 38 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 38, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24964
  Nexthop addr: 4.4.4.4      cross connect ix: 11, op code: Push

Primary FTN entry with FEC: 30.1.1.0/24, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 00:32:22, UpTime: 00:32:22, LastUpdate: 00:11:33
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 5 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 5, owner: LDP, Stale: NO, refcount: 2, out intf: ce6/2, out label: 24963
  Nexthop addr: 10.1.1.2      cross connect ix: 3, op code: Push

  Backup Cross connect ix: 13, in intf: - in label: 0 out-segment ix: 41 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 41, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24967
  Nexthop addr: 60.1.1.2      cross connect ix: 13, op code: Push

Primary FTN entry with FEC: 40.1.1.0/24, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 00:11:33, UpTime: 00:11:33, LastUpdate: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 44 refcount: 1
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 44, owner: LDP, Stale: NO, refcount: 2, out intf: xe29, out label: 24968
  Nexthop addr: 60.1.1.2      cross connect ix: 7, op code: Push

  Backup Cross connect ix: 15, in intf: - in label: 0 out-segment ix: 45 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 45, owner: LDP, Stale: NO, refcount: 2, out intf: ce6/2, out label: 24965
  Nexthop addr: 10.1.1.2      cross connect ix: 15, op code: Push

Primary FTN entry with FEC: 50.1.1.0/24, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
CreateTime: 00:11:33, UpTime: 00:11:33, LastUpdate: N/A
```

```

Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
Algorithm Number:0
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 33 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 33, owner: N/A, Stale: NO, refcount: 4, out intf: xe29, out label: 3
  Nexthop addr: 60.1.1.2      cross connect ix: 6, op code: Push

Backup Cross connect ix: 17, in intf: - in label: 0 out-segment ix: 49 bypass ftn-ix: 0
  Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 49, owner: LDP, Stale: NO, refcount: 1, out intf: ce6/2, out label: 24967
  Nexthop addr: 4.4.4.4      cross connect ix: 17, op code: Push

```

Verify ILM entries in ILM table

```

PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
S - SNMP, L - LDP, R - RSVP, C - CRLDP
B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
P - SR Policy, U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled
Code   FEC/VRF/L2CKT   ILM-ID   In-Label   Out-Label   In-Intf   Out-
Intf/VRF Nexthop      pri  Algo-Num  UpTime      UPStr peers
L>  40.1.1.0/24    4        24965      24968      N/A       xe29      60.1.1.2
      Yes N/A      00:11:49  2
      24965      24965      N/A       ce6/2      10.1.1.2
      No -        -
L>  3.3.3.3/32     7        24962      24962      N/A       ce6/2      10.1.1.2
      Yes N/A      00:11:38  2
      24962      24962      N/A       xe29      4.4.4.4
      No -        -
      (via
60.1.1.2, label 24963)
L>  2.2.2.2/32     6        24960      3          N/A       ce6/2      10.1.1.2
      Yes N/A      00:11:38  2
      24960      24961      N/A       xe29      4.4.4.4
      No -        -
      (via
60.1.1.2, label 24963)
L>  20.1.1.0/24    10       24961      3          N/A       ce6/2      10.1.1.2
      Yes N/A      00:11:38  2
      24961      24964      N/A       xe29      4.4.4.4
      No -        -
      (via
60.1.1.2, label 24963)
L>  4.4.4.4/32     8        24964      24964      N/A       ce6/2      10.1.1.2
      Yes N/A      00:11:38  1
      24964      24963      N/A       xe29      60.1.1.2
      No -        -
      24964      24963      N/A       xe29      60.1.1.2
      Yes -        -
      24964      24964      N/A       ce6/2      10.1.1.2
      No -        -
L>  30.1.1.0/24    11       24963      24963      N/A       ce6/2      10.1.1.2
      Yes N/A      00:11:38  2
      24963      24967      N/A       xe29      60.1.1.2
      No -        -
L>  1.1.1.1/32     5        24969      Nolabel    N/A       N/A        127.0.0.1
      Yes N/A      00:11:38  1
L>  50.1.1.0/24    2        24967      3          N/A       xe29      60.1.1.2
      Yes N/A      00:11:54  2
L>  5.5.5.5/32     3        24966      24964      N/A       xe29      60.1.1.2
      Yes N/A      00:11:49  2
L>  6.6.6.6/32     1        24968      3          N/A       xe29      60.1.1.2
      Yes N/A      00:11:54  2
L>  10.1.1.0/24    9        24970      Nolabel    N/A       N/A        127.0.0.1

```

| | | | | | | | | | | | |
|----|-------------|-----|-----|----------|---|----|-------|---------|-----|-----|-----------|
| L> | 60.1.1.0/24 | Yes | N/A | 00:11:38 | 1 | 12 | 24971 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | Yes | N/A | 00:11:38 | 1 | | | | | | |

**Notes:**

The following CLI will be used to configure different types of MD5 authentication:

- To configure dedicated MD5 password to a neighbor (under router LDP), use the command `neighbor A.B.C.D auth md5 password (plain-text|encrypt) WORD`. The same should be configured on neighbour A.B.C.D.
- To set a password for all LDP neighbors (under router LDP), use the command `neighbor all auth md5 password (plain-text|encrypt) WORD`.
- To exclude a password for a neighbor (under router LDP), use the command [neighbor auth md5 password \(page 572\)](#).
- To set a password for auto-targeted sessions (under router LDP), use the command `neighbor auto-targeted auth md5 password (plain-text|encrypt) WORD`.
- To create a session group (under router LDP), use the command [session-group name \(page 583\)](#).
- To set a password for the session group (under session group), use the command [auth md5 password \(page 518\)](#).
- To add a neighbors in the group (under session group), use the command [session-group name \(page 583\)](#).

| | |
|--|---|
| mode | |
| (config-router)#transport-address ipv4 111.111.111.111 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface xe13 | Enter interface mode |
| (config-if)#speed 1g | Configure interface speed to 1g |
| (config-if)#ip address 10.0.1.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po1000 | Enter the Interface mode for po1000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface po1000.1000 | Enter interface mode for po1000.1000 |
| (config-if)#encapsulation dot1q 1000 | Configure encapsulation under a subinterface |
| (config-if)#ip address 10.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe4 | Enter interface mode |
| (config-if)#speed 1g | Configure interface speed to 1g |
| (config-if)#ip address 16.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe4.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 16.0.1.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |

| | |
|---|---|
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface xe12 | Enter interface mode |
| (config-if)#speed 1g | Configure interface speed to 1g |
| (config-if)#channel-group 1000 mode active | Moving interface to Dynamic LAG |
| (config-if)#exit | Exit interface mode |
| (config)#ip vrf l3vpnvrf300 | IP VRF config with name l3vpnvrf300 |
| (config-vrf)#rd 300:1 | Route-distinguisher value |
| (config-vrf)#route-target both 300:1 | Route target value |
| (config-vrf)#exit | Exit to config mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 111.111.111.111 | Configure ospf Router-id |
| (config-router)#network 10.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 16.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 16.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 111.111.111.111/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#router bgp 100 | Configure router bgp in AS 100 |
| (config-router)#bgp router-id 111.111.111.111 | Configure BGP router ID |
| (config-router)#neighbor 222.222.222.222 remote-as 100 | Configure neighbor in remote-as 100 |
| (config-router)#neighbor 222.222.222.222 update-source lo | Configure neighbor with update-source loopback |
| (config-router)#address-family vpnv4 unicast | Enter VPNv4 Address family mode |
| (config-router-af)#neighbor 222.222.222.222 activate | mode Activate VPNv4 neighbor |
| (config-router-af)#exit-address-family | Exit from address Fmily configuration |
| (config-router)#address-family l2vpn vpls | Enter VPLS address family mode |
| (config-router-af)#neighbor 222.222.222.222 activate | mode Activate vpls neighbor |

| | |
|--|---|
| (config-router-af) #exit-address-family | Exit from Address Family configuration |
| (config-router) #address-family vrf l3vpnvr300 | Configure VRF address family |
| (config-router-af) #redistribute connected | Redistribute connected addresses |
| (config-router-af) #exit-address-family | Exit from Address Family configuration |
| (config-router) #exit | Exit from router mode |
| (config) #mpls vplsldp100 100 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls) #signaling ldp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig) # vpls-peer 222.222.222.222 | Configuring VPLS mesh peers |
| (config-vpls-sig) #exit-signaling | Exit from VPLS signaling mode |
| (config-vpls) #exit-vpls | Exit from VPLS mode |
| (config) #mpls l2-circuit VPWS400 400 222.222.222.222 | Configuring VPWS instance with name and VPWS ID |
| (config) #mpls vpls vplsbgp200 200 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls) #signaling bgp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig) #ve-id 200 | Configure VE ID, which is mandatory for BGP VPLS, otherwise, signaling does not take place. VE ID should be unique per VPLS instance. |
| (config-vpls-sig) #exit-signaling | Exit from VPLS signaling mode |
| (config-vpls) #exit-vpls | Exit from VPLS mode |
| (config-if) #interface xe8.100 switchport | Enter sub interface mode |
| (config-if) #encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if) #access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls) #mpls-vpls vplsldp100 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls) #exit | Exit from access mode |
| (config-if) #interface xe8.400 switchport | Enter sub interface mode |
| (config-if) #encapsulation dot1q 400 | Configure encapsulation under a subinterface |
| (config-if) #access-if-vpws | Access VPWS under sub interface |
| (config-acc-if-vpws) #mpls-vpws VPWS400 | Associating the VPWS Instance to the attachment circuit interface. |
| (config-acc-if-vpws) #exit | Exit from access mode |
| (config-if) #interface xe8.200 switchport | Enter sub interface mode |
| (config-if) #split-horizon group access1 | Configure split-horizon group on sub-interface |
| (config-if) #split-horizon group access1 | Configure split-horizon group on sub-interface |
| (config-if) #encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if) #access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls) #mpls-vpls vplsbgp200 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls) #exit | Exit from access mode |

| | |
|---|--|
| (config-if)#interface xe8.300 | Enter sub interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#ip vrf forwarding l3vpnvrf300 | Attaching xe8.300 to as part of l3vpnvrf300 |
| (config-if)#ip address 110.110.110.1/24 | Configure the IP address of the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the configuration |
| (config)#end | Return to privilege mode |

P1

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#mpls ilm-ecmp | Enable ILM ECMP |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 11.11.11.11/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 11.11.11.11 | Configure Router-id |
| (config-router)#transport-address ipv4 11.11.11.11 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface ge13 | Enter interface mode |
| (config-if)#ip address 10.0.1.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po1000 | Enter the Interface mode for po1000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface po1000.1000 | Enter interface mode for po1000.1000 |
| (config-if)#encapsulation dot1q 1000 | Configure encapsulation under a subinterface |
| (config-if)#ip address 10.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |

| | |
|--|---|
| (config-if)#exit | Exit interface mode |
| (config)#interface po2000 | Enter the Interface mode for po2000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface po2000.100 | Enter interface mode for po2000.100 |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#ip address 11.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge4 | Enter interface mode |
| (config-if)#ip address 9.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 20 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge12 | Enter interface mode |
| (config-if)#channel-group 1000 mode active | Moving interface to Dynamic LAG |
| (config)#interface xe23 | Enter interface mode |
| (config-if)#channel-group 2000 mode active | Moving interface to Dynamic LAG |
| (config-if)#speed 1g | Configure speed 1g |
| (config)#interface xe24 | Enter interface mode |
| (config-if)#channel-group 2000 mode active | Moving interface to Dynamic LAG |
| (config-if)#speed 1g | Configure speed 1g |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 11.11.11.11 | Configure ospf Router-id |
| (config-router)#network 10.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 11.0.0.0/24 area 0 | Define the interface on which OSPF runs and |

| | |
|---|---|
| | associate the area ID (0) with the interface. |
| (config-router)#network 9.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 11.11.11.11/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#exit | Exit from router mode |
| (config)#commit | Commit the transaction |

P2

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 22.22.22.22/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 22.22.22.22 | Configure Router ID |
| (config-router)#transport-address ipv4 22.22.22.22 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface ge6.300 | Enter interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#ip address 15.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge6.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 15.0.1.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po2000 | Enter the Interface mode for po2000 |

| | |
|---|---|
| (config-if)#exit | Exit interface mode |
| (config)#interface po2000.100 | Enter interface mode for po2000.100 |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#ip address 11.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge5 | Enter interface mode |
| (config-if)#channel-group 2000 mode active | Moving interface to Dynamic LAG |
| (config)#interface ge11 | Enter interface mode |
| (config-if)#channel-group 2000 mode active | Moving interface to Dynamic LAG |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 22.22.22.22 | Configure ospf Router ID |
| (config-router)#network 15.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 15.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 11.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 22.22.22.22/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#exit | Exit from router mode |
| (config)#commit | Commit the transaction |

P3

| | |
|---|---|
| #configure terminal | Enter Configure mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 33.33.33.33/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |

| | |
|--|---|
| (config-router)#router-id 33.33.33.33 | Configure Router-id |
| (config-router)#transport-address ipv4 33.33.33.33 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface ge5 | Enter interface mode |
| (config-if)#ip address 16.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge5.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 16.0.1.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa4000 | Enter the Interface mode for sa4000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa4000.200 | Enter interface mode for psa4000.200 |
| (config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if)#ip address 17.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe12 | Enter interface mode |
| (config-if)# static-channel-group 4000 | Moving interface to Static LAG |
| (config)#interface xe13 | Enter interface mode |
| (config-if)# static-channel-group 4000 | Moving interface to Static LAG |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the |

| | |
|---|---|
| | Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 33.33.33.33 | Configure ospf Router-id |
| (config-router)#network 16.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 16.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 17.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 33.33.33.33/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#exit | Exit from router mode |
| (config)#commit | Commit the transaction. |

P4

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 44.44.44.44/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 44.44.44.44 | Configure Router ID |
| (config-router)#transport-address ipv4 44.44.44.44 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface sa3000 | Enter the Interface mode for sa3000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa3000 | Enter interface mode |
| (config-if)#ip address 18.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa3000.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |

| | |
|--|---|
| (config-if)#ip address 18.0.1.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa4000 | Enter the Interface mode for sa4000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa4000.200 | Enter interface mode for psa4000.200 |
| (config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if)#ip address 17.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe6 | Enter interface mode |
| config-if)#speed 1g | Configure speed 1g |
| (config-if)#static-channel-group 3000 | Moving interface to Static LAG |
| (config)#interface xe8 | Enter interface mode |
| config-if)#speed 1g | Configure speed 1g |
| (config-if)#static-channel-group 3000 | Moving interface to Static LAG |
| (config)#interface xe12 | Enter interface mode |
| (config-if)# static-channel-group 4000 | Moving interface to Static LAG |
| (config)#interface xe13 | Enter interface mode |
| (config-if)# static-channel-group 4000 | Moving interface to Static LAG |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 44.44.44.44 | Configure OSPF Router ID |
| (config-router)#network 18.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 18.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 17.0.0.0/24 area 0 | Define the interface on which OSPF runs and |

| | |
|---|---|
| | associate the area ID (0) with the interface. |
| (config-router)#network 44.44.44.44/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#exit | Exit from router mode |
| (config)#commit | Commit the transaction. |

PE2

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 222.222.222.222/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 222.222.222.222 | Configure Router-id |
| (config-router)#targeted-peer ipv4 111.111.111.111 | Configuring targeted LDP sessions to PE-2 |
| (config-router)#explicit-null | Configure explicit-null. |
| (config-router)#entropy-label-capability | Enable entropy capability in ldp |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)#transport-address ipv4 222.222.222.222 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface xe14 | Enter interface mode |
| (config-if)#speed 1g | Configure interface speed to 1g |
| (config)#interface xe14.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 15.0.1.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe14.1002 | Enter interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#ip address 15.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |

| | |
|--|---|
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa3000 | Enter interface mode |
| (config-if)#ip address 18.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa3000.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 18.0.1.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface ge5 | Enter interface mode |
| (config-if)# static-channel-group 3000 | Moving interface to static LAG |
| (config-if)#interface ge7 | Enter interface mode |
| (config-if)# static-channel-group 3000 | Moving interface to static LAG |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge9 | Enter interface mode |
| (config-if)#ip address 9.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 20 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#ip vrf l3vpnvrf300 | IP VRF config with name l3vpnvrf300 |
| (config-vrf)#rd 300:1 | Route-distinguisher value |
| (config-vrf)#route-target both 300:1 | Route target value |
| (config-vrf)#exit | Exit to config mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a |

| | |
|---|---|
| | unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 222.222.222.222 | Configure ospf Router-id |
| (config-router)#network 9.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 15.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 15.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 18.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 18.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 222.222.222.222/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#router bgp 100 | Configure router bgp in AS 100 |
| (config-router)#bgp router-id 222.222.222.222 | Configure BGP router ID |
| (config-router)#neighbor 111.111.111.111 remote-as 100 | Configure neighbor in remote-as 100 |
| (config-router)#neighbor 111.111.111.111 update-source lo | Configure neighbor with update-source loopback |
| (config-router)#address-family vpnv4 unicast | Enter VPNv4 Address family mode |
| (config-router-af)#neighbor 111.111.111.111 activate | mode Activate VPNv4 neighbor |
| (config-router-af)#exit-address-family | Exit from Address Family configuration |
| (config-router)#address-family l2vpn vpls | Enter VPLS Address family mode |
| (config-router-af)#neighbor 111.111.111.111 activate | Activate VPLS neighbor |
| (config-router-af)#exit-address-family | Exit from Address Family configuration |
| (config-router)#address-family vrf l3vpnvr300 | Configure VRF address family |
| (config-router-af)#redistribute connected | Redistribute connected addresses |
| (config-router-af)#exit-address-family | Exit from Address Family configuration |
| (config-router)#exit | Exit from router mode |
| (config)#mpls vplsldp100 100 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig)# vpls-peer 111.111.111.111 | Configuring VPLS mesh peers |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |

| | |
|---|--|
| (config)#mpls l2-circuit VPWS400 400 111.111.111.111 | Configuring VPWS instance with name and VPWS ID |
| (config)#mpls vpls vplsbgp200 200 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls)#signaling bgp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig)#ve-id 201 | Configure VE ID, which is mandatory for BGP VPLS, otherwise, signaling does not take place. VE ID should be unique per VPLS instance |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |
| (config-if)#interface xe12.100 switchport | Enter sub interface mode |
| (config-if)#split-horizon group access1 | Configure split-horizon group on sub-interface |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls vplsldp100 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode |
| (config-if)#interface xe12.400 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 400 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpws | Access VPWS under sub interface |
| (config-acc-if-vpws)#mpls-vpws VPWS400 | Associating the VPWS Instance to the attachment circuit interface. |
| (config-acc-if-vpws)#exit | Exit from access mode |
| (config-if)#interface xe12.200 switchport | Enter sub interface mode |
| (config-if)#split-horizon group access1 | Configure split-horizon group on sub-interface |
| (config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls vplsbgp200 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode |
| (config-if)#interface xe12.300 | Enter sub interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#ip vrf forwarding l3vpnvr300 | Attaching xe12.300 to as part of l3vpnvr300 |
| (config-if)#ip address 210.210.210.1/24 | Configure the IP address of the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the configuration |
| (config)#end | Return to privilege mode |

Validation

PE1

```

---
---
PE1#show ip ospf neighbor

Total number of full neighbors: 4
OSPF process 100 VRF(default):
Neighbor ID      Pri   State           Dead Time   Address      Interface     Instance ID
11.11.11.11      1     Full/Backup     00:00:30    10.0.0.20    po1000.1000   0
11.11.11.11      1     Full/Backup     00:00:34    10.0.1.20    xe13          0
33.33.33.33      1     Full/Backup     00:00:34    16.0.0.20    xe4           0
33.33.33.33      1     Full/Backup     00:00:34    16.0.1.20    xe4.1001      0
PE1#

PE1#show ip ospf interface brief
Interface      PID   Area           Intf ID      Cost   State         Neighbors   Status
lo             100   0.0.0.0        1            1     Loopback     0           Up

Interface      PID   Area           Intf ID      Cost   State         Neighbors   Status
xe4            100   0.0.0.0        10005        10     DR            1           Up

Interface      PID   Area           Intf ID      Cost   State         Neighbors   Status
xe13           100   0.0.0.0        10014        10     DR            1           Up

Interface      PID   Area           Intf ID      Cost   State         Neighbors   Status
xe4.1001       100   0.0.0.0        327844841    10     DR            1           Up

Interface      PID   Area           Intf ID      Cost   State         Neighbors   Status
po1000.1000    100   0.0.0.0        524289000    10     DR            1           Up
PE1#

PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name      My Role      State          KeepAlive  UpTime
11.11.11.11            xe13         Active       OPERATIONAL   30          00:00:33
33.33.33.33            xe4.1001     Active       OPERATIONAL   30          00:00:33
222.222.222.222        xe4.1001     Passive      OPERATIONAL   30          00:00:22

PE1#show bgp neighbors
BGP neighbor is 222.222.222.222, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 111.111.111.111, remote router ID 222.222.222.222
  BGP state = Established, up for 00:00:27
  Last read 00:00:08, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family L2VPN VPLS: advertised and received
    Address family L2VPN EVPN: advertised and received
    Address family IPv6 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Address family IPv6 Labeled Unicast: advertised and received
  Received 11 messages, 0 notifications, 0 in queue
  Sent 13 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo

```

```

For address family: VPNv4 Unicast
  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  1 announced prefixes

```

```

For address family: L2VPN VPLS
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  1 announced prefixes

```

```

Connections established 1; dropped 0
Local host: 111.111.111.111, Local port: 179
Foreign host: 222.222.222.222, Foreign port: 35033
Nextthop: 111.111.111.111
Nextthop global: ::
Nextthop local: ::
BGP connection: non shared network

```

```
PE1#
```

```
PE1#show mpls vpls
```

| Name | VPLS-ID | Type | MPeers | SPeers | SIG- |
|-------------------|---------|----------|--------|--------|------|
| Protocol Learning | | | | | |
| vplsldp100 | 100 | Ethernet | 1 | 0 | LDP |
| nabled | | | | | |
| vplsbgp200 | 200 | Ethernet | 1 | 0 | BGP |
| nabled | | | | | |

```
PE1#
```

```
PE1#show mpls vpls mesh
```

```

(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 100 | 222.222.222.222 | 25611 | 26253 | xe13 | 26246 | 2/Up | 1 L |
| DP | Active | 00:56:04 | - | | | | |
| 200 | 222.222.222.222 | 25611 | 25608 | xe13 | 25607 | 2/Up | 3 BGP |
| | Active | 00:56:10 | - | | | | |

```
PE1#show mpls vc-table
```

```

(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VC-ID | Access-Intf | Network-Intf | In label | Out Label | Tunnel- |
|-------|-----------------|--------------|----------|-----------|---------|
| Label | Nextthop | Status | UpTime | | |
| 400 | xe8.400 | - | 26290 | 26252 | N/A |
| (e) | 222.222.222.222 | Active | - | | |

```
PE1#show mpls l2-circuit
```

```

MPLS Layer-2 Virtual Circuit: VPWS400, id: 400 PW-INDEX: 2 service-tpid: dot1q
Endpoint: 222.222.222.222
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe8.400
Subinterface Match Criteria(s) :
dot1q 400

```

```

Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active

```

```
PE1#
```

```
PE1#show mpls vrf-forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, B - BGP FTN
(m) - Service mapped over multipath transport

```

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|------|------------------|--------|-----------|-----|-------------|-----------|------|
| Intf | Nexthop | | | | | | |
| B> | 210.210.210.0/24 | 1 | - | - | LSP_DEFAULT | 25664 | - |
| | 222.222.222.222 | | | | | | |
| B> | 210::/64 | 2 | - | - | LSP_DEFAULT | 25664 | - |
| | 222.222.222.222 | | | | | | |

```
PE1#
```

```
PE1#show mpls forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

| Code | FEC | FTN-ID | Nhife-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|-------------------------|----------------|--------|-----------------|-----------|-----|----------|-----------|------|
| Intf | ELC Nexthop | | | | | | | |
| L> | 9.0.0.0/24 | 1 | 73(e) | | | | | |
| | | | 52 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe13 | No | 10.0.1.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | (via |
| | | | | | | | | |
| 16.0.0.20 ,label 25612) | | | | | | | | |
| | | | 53 | - | Yes | LSP_ | | |
| DEFAULT | 3 po1000.1000 | No | 10.0.0.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | (via |
| | | | | | | | | |
| 16.0.0.20 ,label 25612) | | | | | | | | |
| L> | 11.0.0.0/24 | 2 | 73(e) | | | | | |
| | | | 52 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe13 | No | 10.0.1.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | (via |
| | | | | | | | | |
| 16.0.0.20 ,label 25612) | | | | | | | | |
| | | | 53 | - | Yes | LSP_ | | |
| DEFAULT | 3 po1000.1000 | No | 10.0.0.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | (via |
| | | | | | | | | |
| 16.0.0.20 ,label 25612) | | | | | | | | |
| L> | 11.11.11.11/32 | 3 | 100(e) | | | | | |
| | | | 52 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe13 | No | 10.0.1.20 | | | | | |
| | | | 53 | - | No | LSP_ | | |
| DEFAULT | 3 po1000.1000 | No | 10.0.0.20 | | | | | |
| | | | 53 | - | Yes | LSP_ | | |
| DEFAULT | 3 po1000.1000 | No | 10.0.0.20 | | | | | |
| | | | 52 | - | No | LSP_ | | |
| DEFAULT | 3 xe13 | No | 10.0.1.20 | | | | | |
| L> | 15.0.0.0/24 | 4 | 77(e) | | | | | |
| | | | 56 | - | Yes | LSP_ | | |
| DEFAULT | 25600 xe13 | No | 10.0.1.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | |


```

16.0.0.20 ,label 25612)                                     (via
                    57      -      Yes  LSP_
DEFAULT 25600      po1000.1000 No 10.0.0.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No 222.222.222.222

16.0.0.20 ,label 25612)                                     (via
L> 15.0.1.0/24      5          80(e)
                    60      -      Yes  LSP_
DEFAULT 25601      xe13        No 10.0.1.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No 222.222.222.222

16.0.0.20 ,label 25612)                                     (via
                    61      -      Yes  LSP_
DEFAULT 25601      po1000.1000 No 10.0.0.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No 222.222.222.222

16.0.0.20 ,label 25612)                                     (via
L> 17.0.0.0/24      6          84(e)
                    16      -      Yes  LSP_
DEFAULT 3          xe4          No 16.0.0.20
                    82      -      No   LSP_
DEFAULT 26249      po1000.1000 No 222.222.222.222

10.0.0.20 ,label 25611)                                     (via
                    17      -      Yes  LSP_
DEFAULT 3          xe4.1001    No 16.0.1.20
                    82      -      No   LSP_
DEFAULT 26249      po1000.1000 No 222.222.222.222

10.0.0.20 ,label 25611)                                     (via
L> 18.0.0.0/24      7          88(e)
                    19      -      Yes  LSP_
DEFAULT 25608      xe4          No 16.0.0.20
                    86      -      No   LSP_
DEFAULT 0          po1000.1000 No 222.222.222.222

10.0.0.20 ,label 25611)                                     (via
                    20      -      Yes  LSP_
DEFAULT 25608      xe4.1001    No 16.0.1.20
                    86      -      No   LSP_
DEFAULT 0          po1000.1000 No 222.222.222.222

10.0.0.20 ,label 25611)                                     (via
L> 18.0.1.0/24      8          91(e)
                    22      -      Yes  LSP_
DEFAULT 25609      xe4          No 16.0.0.20
                    86      -      No   LSP_
DEFAULT 0          po1000.1000 No 222.222.222.222

10.0.0.20 ,label 25611)                                     (via
                    23      -      Yes  LSP_
DEFAULT 25609      xe4.1001    No 16.0.1.20
                    86      -      No   LSP_
DEFAULT 0          po1000.1000 No 222.222.222.222

10.0.0.20 ,label 25611)                                     (via
L> 22.22.22.22/32   9          94(e)
                    64      -      Yes  LSP_
DEFAULT 25607      xe13        No 10.0.1.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No 222.222.222.222

16.0.0.20 ,label 25612)                                     (via
                    65      -      Yes  LSP_
DEFAULT 25607      po1000.1000 No 10.0.0.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No 222.222.222.222

```

(via

```
16.0.0.20 ,label 25612)
L> 33.33.33.33/32 10 102(e)
16 - Yes LSP_
DEFAULT 3 xe4 No 16.0.0.20
17 - No LSP_
DEFAULT 3 xe4.1001 No 16.0.1.20
17 - Yes LSP_
DEFAULT 3 xe4.1001 No 16.0.1.20
16 - No LSP_
DEFAULT 3 xe4 No 16.0.0.20
L> 44.44.44.44/32 11 98(e)
28 - Yes LSP_
DEFAULT 25611 xe4 No 16.0.0.20
96 - No LSP_
DEFAULT 26251 po1000.1000 No 222.222.222.222

(via
10.0.0.20 ,label 25611)
29 - Yes LSP_
DEFAULT 25611 xe4.1001 No 16.0.1.20
96 - No LSP_
DEFAULT 26251 po1000.1000 No 222.222.222.222

(via
10.0.0.20 ,label 25611)
L> 222.222.222.222/32 12 104(e)
68 - Yes LSP_
DEFAULT 25611 xe13 Yes 10.0.1.20
31 - No LSP_
DEFAULT 25612 xe4 Yes 16.0.0.20
69 - Yes LSP_
DEFAULT 25611 po1000.1000 Yes 10.0.0.20
31 - No LSP_
DEFAULT 25612 xe4 Yes 16.0.0.20
31 - Yes LSP_
DEFAULT 25612 xe4 Yes 16.0.0.20
69 - No LSP_
DEFAULT 25611 po1000.1000 Yes 10.0.0.20
32 - Yes LSP_
DEFAULT 25612 xe4.1001 Yes 16.0.1.20
69 - No LSP_
DEFAULT 25611 po1000.1000 Yes 10.0.0.20
B> 220::/64 13 49 - LSP_DEFAULT 25665 -
No 222.222.222.222

PE1#
```

```
PE1#show interface counters rate mbps
```

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-------------|---------|--------|---------|--------|
| po1000 | 350.55 | 4850 | 250.40 | 3466 |
| po1000.1000 | 350.55 | 4851 | 250.39 | 3465 |
| xe0 | 0.00 | 0 | 0.00 | 0 |
| xe2 | 0.00 | 0 | 0.00 | 0 |
| xe4 | 400.63 | 5546 | 500.68 | 6932 |
| xe4.1001 | 400.63 | 5545 | 250.40 | 3465 |
| xe8 | 997.81 | 13858 | 997.81 | 13858 |
| xe8.400 | 997.80 | 13858 | 997.80 | 13858 |
| xe12 | 350.56 | 4850 | 250.40 | 3466 |
| xe13 | 250.28 | 3467 | 250.29 | 3466 |

```
PE1#
```

P1

```
---
```

```
P1#show ip ospf interface brief
```

| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
|-----------|-----|------|---------|------|-------|-----------|--------|
|-----------|-----|------|---------|------|-------|-----------|--------|

```

lo          100  0.0.0.0      1          1      Loopback      0          Up

Interface   PID   Area          Intf ID    Cost   State          Neighbors  Status
ge8         100  0.0.0.0      10009     20    Backup         1          Up

Interface   PID   Area          Intf ID    Cost   State          Neighbors  Status
ge13        100  0.0.0.0      10014     10    Backup         1          Up

Interface   PID   Area          Intf ID    Cost   State          Neighbors  Status
po1000.1000 100  0.0.0.0      524289000 10    Backup         1          Up

Interface   PID   Area          Intf ID    Cost   State          Neighbors  Status
po2000.100  100  0.0.0.0      557056100 10    Backup         1          Up

P1#

P1#show ip ospf neighbor

Total number of full neighbors: 4
OSPF process 100 VRF(default):
Neighbor ID  Pri  State          Dead Time  Address        Interface      Instance ID
222.222.222.222  1  Full/DR        00:00:38   9.0.0.20      ge8            0
111.111.111.111  1  Full/DR        00:00:36   10.0.0.10     po1000.1000    0
111.111.111.111  1  Full/DR        00:00:34   10.0.1.10     ge13           0
22.22.22.22      1  Full/DR        00:00:38   11.0.0.20     po2000.100     0
P1#

P2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name      My Role      State          KeepAlive  UpTime
222.222.222.222      ge6.1001     Passive      OPERATIONAL   30          00:00:21
11.11.11.11          po2000.100   Active       OPERATIONAL   30          00:01:27
33.33.33.33          po2000.100   Passive      OPERATIONAL   30          00:00:14

P1#show interface counters rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
ge8         190.14   2634    250.19   3464
ge10        0.00     0        0.00     0
ge12        250.30   3465    350.42   4848
ge13        250.19   3465    250.19   3465
po1000       250.30   3465    350.42   4849
po1000.1000  250.31   3465    350.42   4848
po2000       410.50   5680    250.30   3464
po2000.100  410.50   5681    250.30   3464
xe23        220.26   3048    120.15   1663
xe24        190.23   2632    130.16   1801
xe26        0.00     0        0.00     0
P1#

P1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

LDP ilm-ecmp - enabled
Code  FEC/VRF/L2CKT  ILM-ID      In-Label    Out-Label    In-Intf      Out-
Intf/VRF      Nexthop     pri         LSP-Type    UpTime       UPStr-peers

L> 111.111.111.111/32

```

| | | | | | | | |
|---------|------------------------|-------------|----------|-------|-----------|-------------|------------|
| 22 | 25621 | 0 | N/A | ge13 | 10.0.1.10 | Yes | LSP_ |
| DEFAULT | 00:04:23 | 1 | | | | | |
| | No | LSP_DEFAULT | 25621 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | | | 00:04:23 | | | | |
| | Yes | LSP_DEFAULT | 25621 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | | | 00:04:23 | | | | |
| | No | LSP_DEFAULT | 25621 | 0 | N/A | ge13 | 10.0.1.10 |
| | | | 00:04:23 | | | | |
| L> | 16.0.1.0/24 | 9 | 25608 | 0 | N/A | ge13 | 10.0.1.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25608 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 15.0.0.0/24 | 3 | 25602 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25602 | 0 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 15.0.1.0/24 | 2 | 25601 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25601 | 0 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 15.0.0.0/24 | 1 | 25600 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25600 | 0 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | 1 | | | |
| L> | 16.0.0.0/24 | 5 | 25604 | 0 | N/A | ge13 | 10.0.1.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25604 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 15.0.1.0/24 | 4 | 25603 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25603 | 0 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 18.0.1.0/24 | 7 | 25606 | 25604 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25606 | 0 | N/A | ge8 | 9.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25606 | 3 | N/A | ge8 | 44.44.44.4 |
| 4 | No | LSP_DEFAULT | 00:04:23 | | | | |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 18.0.0.0/24 | 6 | 25605 | 25603 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25605 | 0 | N/A | ge8 | 9.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25605 | 3 | N/A | ge8 | 44.44.44.4 |
| 4 | No | LSP_DEFAULT | 00:04:23 | | | | |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 22.22.22.22/32 | 8 | 25607 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25607 | 26267 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 111.111.111.111/32 | | | | | | |
| 17 | 25616 | 0 | N/A | ge13 | 10.0.1.10 | Yes | LSP_ |
| DEFAULT | 00:04:23 | | | | | | |
| | No | LSP_DEFAULT | 25616 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | | | 00:04:23 | | | | |
| | Yes | LSP_DEFAULT | 25616 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | | | 00:04:23 | | | | |
| | No | LSP_DEFAULT | 25616 | 0 | N/A | ge13 | 10.0.1.10 |
| | | | 00:04:23 | | | | |
| L> | 16.0.1.0/24 | 14 | 25613 | 0 | N/A | ge13 | 10.0.1.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25613 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 18.0.0.0/24 | 11 | 25610 | 25603 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25610 | 0 | N/A | ge8 | 9.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | 1 | | | |
| L> | 17.0.0.0/24 | 10 | 25609 | 26245 | N/A | ge13 | 10.0.1.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |

| | | | | | | | |
|---------|------------------------|-----------------|----------|------------|-----------|-------------|-------------|
| | | | 25609 | 26245 | N/A | po1000.1000 | 10.0.0.10 |
| L> | 16.0.0.0/24 | Yes LSP_DEFAULT | 00:04:23 | 25612 | 0 | N/A | ge13 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25612 | 0 | N/A | po1000.1000 |
| L> | 222.222.222.222/32 | Yes LSP_DEFAULT | 00:04:23 | | | | |
| 12 | 25611 | 25607 | N/A | po2000.100 | 11.0.0.20 | Yes | LSP_ |
| DEFAULT | 00:04:23 | | | | | | |
| | | Yes LSP_DEFAULT | 00:04:23 | 25611 | 0 | N/A | ge8 |
| 4 | | No LSP_DEFAULT | 00:04:23 | 25611 | 25624 | N/A | ge8 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 17.0.0.0/24 | 15 | 25614 | 26245 | N/A | ge13 | 10.0.1.10 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25614 | 3 | N/A | ge8 |
| 4 | | No LSP_DEFAULT | 00:04:23 | | | | 44.44.44.4 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| | | Yes LSP_DEFAULT | 00:04:23 | 25614 | 26245 | N/A | po1000.1000 |
| 4 | | No LSP_DEFAULT | 00:04:23 | 25614 | 3 | N/A | ge8 |
| | | | | | | | 44.44.44.4 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 33.33.33.33/32 | 16 | 25615 | 26249 | N/A | ge13 | 10.0.1.10 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25615 | 25622 | N/A | ge8 |
| 4 | | No LSP_DEFAULT | 00:04:23 | | | | 44.44.44.4 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| | | Yes LSP_DEFAULT | 00:04:23 | 25615 | 26249 | N/A | po1000.1000 |
| 4 | | No LSP_DEFAULT | 00:04:23 | 25615 | 25622 | N/A | ge8 |
| | | | | | | | 44.44.44.4 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 22.22.22.22/32 | 19 | 25618 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25618 | 26267 | N/A | ge8 |
| | | No LSP_DEFAULT | 00:04:23 | 25617 | 25604 | N/A | po2000.100 |
| L> | 18.0.1.0/24 | 18 | 25617 | 0 | N/A | ge8 | 11.0.0.20 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25617 | 0 | N/A | ge8 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25619 | 26249 | N/A | po1000.1000 |
| L> | 33.33.33.33/32 | 20 | 25619 | 26249 | N/A | ge13 | 10.0.1.10 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25619 | 26249 | N/A | po1000.1000 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 25606 | N/A | po2000.100 |
| L> | 44.44.44.44/32 | 21 | 25620 | 26269 | N/A | ge8 | 11.0.0.20 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 26250 | N/A | ge13 |
| | | No LSP_DEFAULT | 00:04:23 | 25620 | 26269 | N/A | ge8 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 26250 | N/A | po1000.1000 |
| | | No LSP_DEFAULT | 00:04:23 | 25620 | 26269 | N/A | ge8 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 26269 | N/A | ge8 |
| | | No LSP_DEFAULT | 00:04:23 | 25620 | 26250 | N/A | po1000.1000 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 1 | N/A | po1000.1000 |
| L> | 33.33.33.33/32 | 32 | 25631 | 26249 | N/A | ge13 | 10.0.1.10 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25631 | 26249 | N/A | po1000.1000 |
| | | Yes LSP_DEFAULT | 00:04:23 | | | | |

```

L> 16.0.0.0/24      26      25625      0      N/A      ge13      10.0.1.10
      Yes      LSP_DEFAULT      00:04:23      25625      0      N/A      po1000.1000      10.0.0.10
      Yes      LSP_DEFAULT      00:04:23
L> 15.0.0.0/24      24      25623      3      N/A      po2000.100      11.0.0.20
      Yes      LSP_DEFAULT      00:04:23      25623      0      N/A      ge8      9.0.0.20
      No      LSP_DEFAULT      00:04:23
L> 222.222.222.222/32
23      25622      25607      N/A      po2000.100      11.0.0.20      Yes      LSP_
DEFAULT      00:04:23
      Yes      LSP_DEFAULT      00:04:23      25622      0      N/A      ge8      9.0.0.20
L> 15.0.1.0/24      25      25624      3      N/A      po2000.100      11.0.0.20
      Yes      LSP_DEFAULT      00:04:23      25624      0      N/A      ge8      9.0.0.20
      No      LSP_DEFAULT      00:04:23

```

P2

```

---
P2#show ip ospf interface brief
Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
lo             100      0.0.0.0      1             1      Loopback      0              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
ge6.300        100      0.0.0.0      328171820    10      Backup      1              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
ge6.1001       100      0.0.0.0      328172521    10      Backup      1              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
po2000.100     100      0.0.0.0      557056100    10      DR          1              Up

P2#

P2#show ip ospf neighbor

Total number of full neighbors: 3
OSPF process 100 VRF(default):
Neighbor ID      Pri      State      Dead Time      Address      Interface      Instance ID
11.11.11.11      1      Full/Backup      00:00:32      11.0.0.10    po2000.100      0
222.222.222.222  1      Full/DR          00:00:36      15.0.0.20    ge6.300         0
222.222.222.222  1      Full/DR          00:00:35      15.0.1.20    ge6.1001        0
P2#

P2#show ldp session
Codes: m - MD5 password is not set/unset.
      g - GR configuration not set/unset.
      t - TCP MSS not set/unset.
      Session has to be cleared manually

Code  Peer IP Address      IF Name      My Role      State      KeepAlive      UpTime
222.222.222.222      ge6.1001      Passive      OPERATIONAL      30          00:00:21
11.11.11.11          po2000.100    Active      OPERATIONAL      30          00:01:27
33.33.33.33          po2000.100    Passive      OPERATIONAL      30          00:00:14      00:00:14
---
P3#show ip ospf interface brief
Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
lo             100      0.0.0.0      1             1      Loopback      0              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
ge5            100      0.0.0.0      10014         10      Backup      1              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
ge5.1001       100      0.0.0.0      328139753     10      Backup      1              Up

```

```

Interface    PID    Area          Intf ID    Cost    State          Neighbors    Status
sa4000.200   100    0.0.0.0       1179648200 10     Backup         1            Up

P3#

P3#show ip ospf neighbor

Total number of full neighbors: 3
OSPF process 100 VRF(default):
Neighbor ID    Pri    State          Dead Time    Address          Interface          Instance ID
111.111.111.111 1    Full/DR        00:00:38     16.0.0.10       ge5                0
111.111.111.111 1    Full/DR        00:00:32     16.0.1.10       ge5.1001          0
44.44.44.44    1    Full/DR        00:00:36     17.0.0.10       sa4000.200        0
P3#

```

```

P3#show ldp session
Peer IP Address          IF Name    My Role    State          KeepAlive    UpTime
44.44.44.44              sa4000.200 Passive    OPERATIONAL    30           00:02:00
111.111.111.111         ge5        Passive    OPERATIONAL    30           00:00:32
22.22.22.22              ge5        Active     OPERATIONAL    30           00:00:14
11.11.11.11              ge5        Active     OPERATIONAL    30           00:00:20
222.222.222.222         sa4000.200 Passive    OPERATIONAL    30           00:00:13
P3#
P4
---
```

```

P4#show ip ospf interface brief

Interface    PID    Area          Intf ID    Cost    State          Neighbors    Status
lo           100    0.0.0.0       1          1       Loopback      0            Up

Interface    PID    Area          Intf ID    Cost    State          Neighbors    Status
sa3000       100    0.0.0.0       203000     10     Backup        1            Up

Interface    PID    Area          Intf ID    Cost    State          Neighbors    Status
sa3000.1001  100    0.0.0.0       1146881001 10     Backup        1            Up

Interface    PID    Area          Intf ID    Cost    State          Neighbors    Status
sa4000.200   100    0.0.0.0       1179648200 10     DR            1            Up

P4#

```

```

P4#show ip ospf neighbor

Total number of full neighbors: 3
OSPF process 100 VRF(default):
Neighbor ID    Pri    State          Dead Time    Address          Interface          Instance ID
33.33.33.33    1    Full/Backup    00:00:34     17.0.0.20       sa4000.200        0
222.222.222.222 1    Full/DR        00:00:36     18.0.0.20       sa3000            0
222.222.222.222 1    Full/DR        00:00:34     18.0.1.20       sa3000.1001      0
P4#

```

```

P4#show ldp session
Peer IP Address          IF Name    My Role    State          KeepAlive    UpTime
222.222.222.222         sa3000     Passive    OPERATIONAL    30           00:00:40
33.33.33.33             sa4000.200 Active     OPERATIONAL    30           00:02:00
11.11.11.11             sa4000.200 Active     OPERATIONAL    30           00:00:14
P4#

```

PE2

```

---

PE2#show ip ospf neighbor

Total number of full neighbors: 5
OSPF process 100 VRF(default):
Neighbor ID    Pri    State          Dead Time    Address          Interface          Instance ID

```

```

11.11.11.11      1    Full/Backup    00:00:36    9.0.0.10    ge9          0
22.22.22.22      1    Full/Backup    00:00:31    15.0.0.10   xe14.1002    0
22.22.22.22      1    Full/Backup    00:00:34    15.0.1.10   xe14.1001    0
44.44.44.44      1    Full/Backup    00:00:38    18.0.0.10   sa3000       0
44.44.44.44      1    Full/Backup    00:00:31    18.0.1.10   sa3000.1001  0
PE2#

```

PE2#show ip ospf interface brief

| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
|-------------|-----|---------|------------|------|----------|-----------|--------|
| lo | 100 | 0.0.0.0 | 1 | 1 | Loopback | 0 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| ge9 | 100 | 0.0.0.0 | 10018 | 20 | DR | 1 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| sa3000 | 100 | 0.0.0.0 | 203000 | 10 | DR | 1 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| xe14.1001 | 100 | 0.0.0.0 | 328434665 | 10 | DR | 1 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| xe14.1002 | 100 | 0.0.0.0 | 328434666 | 10 | DR | 1 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| sa3000.1001 | 100 | 0.0.0.0 | 1146881001 | 10 | DR | 1 | Up |

PE2#

PE2#show ldp session

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|-----------|---------|-------------|-----------|----------|
| 44.44.44.44 | sa3000 | Active | OPERATIONAL | 30 | 00:00:41 |
| 11.11.11.11 | ge9 | Active | OPERATIONAL | 30 | 00:00:18 |
| 22.22.22.22 | xe14.1002 | Active | OPERATIONAL | 30 | 00:00:25 |
| 111.111.111.111 | xe14.1001 | Active | OPERATIONAL | 30 | 00:00:22 |
| 33.33.33.33 | sa3000 | Active | OPERATIONAL | 30 | 00:00:14 |

PE2#

PE2#show bgp neighbors

```

BGP neighbor is 111.111.111.111, remote AS 100, local AS 100, internal link
BGP version 4, local router ID 222.222.222.222, remote router ID 111.111.111.111
BGP state = Established, up for 00:00:27
Last read 00:00:01, hold time is 90, keepalive interval is 30 seconds
Neighbor capabilities:
  Route refresh: advertised and received (old and new)
  Address family IPv4 Unicast: advertised and received
  Address family VPNv4 Unicast: advertised and received
  Address family L2VPN VPLS: advertised and received
  Address family L2VPN EVPN: advertised and received
  Address family IPv6 Unicast: advertised and received
  Address family VPNv6 Unicast: advertised and received
  Address family IPv6 Labeled Unicast: advertised and received
Received 12 messages, 0 notifications, 0 in queue
Sent 11 messages, 0 notifications, 0 in queue
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 5 seconds
Update source is lo

```

For address family: VPNv4 Unicast

```

BGP table version 2, neighbor version 2
Index 1, Offset 0, Mask 0x2
AIGP is enabled
Community attribute sent to this neighbor (both)
Large Community attribute sent to this neighbor
1 accepted prefixes
1 announced prefixes

```



```

For address family: L2VPN VPLS
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (both)
Large Community attribute sent to this neighbor
0 accepted prefixes
1 announced prefixes

```

```

Connections established 1; dropped 0
Local host: 222.222.222.222, Local port: 35033
Foreign host: 111.111.111.111, Foreign port: 179
Nextthop: 222.222.222.222
Nextthop global: ::
Nextthop local: ::
BGP connection: non shared network

```

PE2#

PE2#show mpls vpls

| Name | VPLS-ID | Type | MPeers | SPeers | SIG- | |
|-------------------|---------|----------|--------|--------|------|---|
| Protocol Learning | | | | | | |
| vplsldp100 | 100 | Ethernet | 1 | 0 | LDP | E |
| nabled | | | | | | |
| vplsbgp200 | 200 | Ethernet | 1 | 0 | BGP | E |
| nabled | | | | | | |

PE2#

PE2#show mpls vpls mesh

(m) - Service mapped over multipath transport

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | |
|--------------|-----------------|--------------|----------|--------------|-----------|---------|----------|---|
| SIG-Protocol | Status | | | | | | | |
| 100 | 111.111.111.111 | 25614 | 26246 | sa3000.1001 | 26253 | 2/Up | 1 | L |
| DP | Active | | | | | | | |
| 200 | 111.111.111.111 | 25614 | 25607 | sa3000.1001 | 25608 | 2/Up | 3 | B |
| GP | Active | | | | | | | |

PE2#

PE2#show mpls vc-table

(m) - Service mapped over multipath transport

(e) - Service mapped over LDP ECMP

| VC-ID | Vlan-ID | Inner-Vlan-ID | Access-Intf | Network-Intf | Out Label | Tunnel- |
|-------|-----------------|---------------|-------------|--------------|-----------|---------|
| Label | Nextthop | Status | | | | |
| 400 | N/A | N/A | xe12.400 | N/A (e) | 26252 | N/A |
| (e) | 111.111.111.111 | Active | | | | |

PE2#

PE2#show mpls l2-circuit

MPLS Layer-2 Virtual Circuit: VPWS400, id: 400 PW-INDEX: 2 service-tpid: dot1q

Endpoint: 111.111.111.111

Control Word: 0

Flow Label Status: Disabled, Direction: None, Static: No

MPLS Layer-2 Virtual Circuit Group: none

Bound to interface: xe12.400

Subinterface Match Criteria(s) :

dot1q 400

Virtual Circuit Type: Ethernet VLAN

Virtual Circuit is configured as Primary

Virtual Circuit is configured as Active

Virtual Circuit is active

PE2#

PE2#show mpls vrf-forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, B - BGP FTN

(m) - Service mapped over multipath transport

| Code Intf | FEC Nexthop | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|--------------|-------------------------------------|--------|-----------|-----|-------------|-----------|------|
| B> | 110.110.110.0/24 111.111.111.111 | 2 | - | - | LSP_DEFAULT | 25664 | - |
| B> | 110::/64 111.111.111.111 | 1 | - | - | LSP_DEFAULT | 25664 | - |

PE2#

PE2#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

| Code Intf | FEC ELC Nexthop | FTN-ID UpTime | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|------------------------|--------------------|------------------|-----------------|-----------|----------|----------|-----------|------|
| L> | 10.0.0.0/24 | 8 | 56(e) | | | | | |
| DEFAULT | 25617 xe14.1001 | No | 54 | - | Yes | LSP_ | | |
| | | | 15.0.1.10 | | 00:03:19 | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 25617 xe14.1002 | No | 55 | - | Yes | LSP_ | | |
| | | | 15.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | Yes | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 0 ge9 | No | 38 | - | No | LSP_ | | |
| | | | 111.111.111.111 | | | | | (via |
| 9.0.0.10 ,label 25642) | | | | | | | | |
| L> | 10.0.1.0/24 | 9 | 59(e) | | | | | |
| DEFAULT | 25618 xe14.1001 | No | 57 | - | Yes | LSP_ | | |
| | | | 15.0.1.10 | | - | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 25618 xe14.1002 | No | 58 | - | Yes | LSP_ | | |
| | | | 15.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | Yes | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 0 ge9 | No | 38 | - | No | LSP_ | | |
| | | | 111.111.111.111 | | | | | (via |
| 9.0.0.10 ,label 25642) | | | | | | | | |
| L> | 11.0.0.0/24 | 11 | 63(e) | | | | | |
| DEFAULT | 3 xe14.1001 | No | 60 | - | Yes | LSP_ | | |
| | | | 15.0.1.10 | | 00:03:19 | | | |
| DEFAULT | 3 xe14.1002 | No | 61 | - | Yes | LSP_ | | |
| | | | 15.0.0.10 | | | | | |
| DEFAULT | 26255 ge9 | No | 62 | - | No | LSP_ | | |
| | | | 111.111.111.111 | | | | | (via |
| 9.0.0.10 ,label 25642) | | | | | | | | |
| L> | 11.11.11.11/32 | 10 | 68(e) | | | | | |
| DEFAULT | 25619 xe14.1001 | No | 66 | - | Yes | LSP_ | | |
| | | | 15.0.1.10 | | - | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 25619 xe14.1002 | No | 67 | - | Yes | LSP_ | | |
| | | | 15.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | Yes | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 40 | - | No | LSP_ | | |

```

DEFAULT 26256 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
L> 16.0.0.0/24 1 42(e)
1 - Yes LSP_
DEFAULT 25604 sa3000.1001 No 18.0.1.10 00:03:19
38 - No LSP_
DEFAULT 0 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
2 - Yes LSP_
DEFAULT 25604 sa3000 No 18.0.0.10
38 - No LSP_
DEFAULT 0 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
L> 16.0.1.0/24 2 43(e)
4 - Yes LSP_
DEFAULT 25605 sa3000.1001 No 18.0.1.10 -
38 - No LSP_
DEFAULT 0 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
5 - Yes LSP_
DEFAULT 25605 sa3000 No 18.0.0.10
38 - No LSP_
DEFAULT 0 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
L> 17.0.0.0/24 3 53(e)
7 - Yes LSP_
DEFAULT 3 sa3000.1001 No 18.0.1.10 00:03:19
52 - No LSP_
DEFAULT 3 ge9 No 33.33.33.33
9.0.0.10 ,label 25641)
8 - Yes LSP_
DEFAULT 3 sa3000 No 18.0.0.10
52 - No LSP_
DEFAULT 3 ge9 No 33.33.33.33
9.0.0.10 ,label 25641)
L> 22.22.22.22/32 12 69(e)
60 - Yes LSP_
DEFAULT 3 xe14.1001 No 15.0.1.10 -
61 - No LSP_
DEFAULT 3 xe14.1002 No 15.0.0.10
61 - Yes LSP_
DEFAULT 3 xe14.1002 No 15.0.0.10
60 - No LSP_
DEFAULT 3 xe14.1001 No 15.0.1.10
L> 33.33.33.33/32 4 45(e)
10 - Yes LSP_
DEFAULT 25606 sa3000.1001 No 18.0.1.10 00:03:19
44 - No LSP_
DEFAULT 26263 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
11 - Yes LSP_
DEFAULT 25606 sa3000 No 18.0.0.10
44 - No LSP_
DEFAULT 26263 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
L> 44.44.44.44/32 5 31(e)
7 - Yes LSP_
DEFAULT 3 sa3000.1001 No 18.0.1.10 -
8 - No LSP_
DEFAULT 3 sa3000 No 18.0.0.10
8 - Yes LSP_
DEFAULT 3 sa3000 No 18.0.0.10

```

```

DEFAULT 3          sa3000.1001 No 7 - No LSP_
L> 111.111.111.111/32 6 73(e) 71 - Yes LSP_
                                36 - 00:03:19
DEFAULT 25622      xe14.1001 Yes 15.0.1.10 No LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 Yes LSP_
                                72 - 15.0.0.10
DEFAULT 25622      xe14.1002 Yes 15.0.0.10 No LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 Yes LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 No LSP_
                                17 - 18.0.0.10
DEFAULT 25614      sa3000 Yes 18.0.0.10 Yes LSP_
                                16 - 18.0.1.10
DEFAULT 25614      sa3000.1001 Yes 18.0.1.10 No LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 Yes LSP_
                                17 - 18.0.0.10
DEFAULT 25614      sa3000 Yes 18.0.0.10 No LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 - LSP_DEFAULT 25665 -
B> 120::/64 7 23 -
    No 111.111.111.111
PE2#

```

PE2#show interface counters rate mbps

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-------------|---------|--------|---------|--------|
| ge4 | 0.00 | 0 | 0.00 | 0 |
| ge5 | 240.28 | 3327 | 210.29 | 2913 |
| ge6 | 0.00 | 0 | 0.00 | 0 |
| ge7 | 260.29 | 3603 | 190.27 | 2635 |
| ge9 | 250.29 | 3466 | 190.22 | 2635 |
| ge10 | 0.00 | 0 | 0.00 | 0 |
| sa3000 | 500.58 | 6930 | 400.56 | 5548 |
| sa3000.1001 | 0.00 | 0 | 200.33 | 2774 |
| xe12 | 997.83 | 13858 | 997.83 | 13858 |
| xe12.400 | 997.83 | 13858 | 997.83 | 13858 |
| xe14 | 250.40 | 3468 | 410.66 | 5685 |
| xe14.1001 | 0.00 | 1 | 58.53 | 811 |
| xe14.1002 | 250.40 | 3467 | 200.33 | 2772 |

PE2#

LDP Graceful Restart


LDP graceful restart (GR) minimizes the negative effects on MPLS traffic caused by control-plane restarts in Label Switching Routers (LSR), especially by the restart of the Label Distribution Protocol (LDP).

LDP graceful restart enables a router whose LDP control plane is undergoing a restart to continue forwarding traffic while recovering its state from neighboring routers. This requires a restarting LDP router that retains established LSP labels. In helper mode, the router maintains label bindings as stale and reprocesses them, after the router undergoing graceful restart reestablishes its LDP session.

The MPLS forwarding state, which is the minimum state required to avoid any disturbance to LSPs traversing a restarting LSR, is preserved during the restart. This mechanism does not require any of the LDP-related states to be preserved across the restart. This means that when LDP restarts, there are minimal or no changes made to the forwarding table entries, and MPLS forwarding continues uninterrupted. This supports Graceful Restart in restarting routers as well as in neighbor routers.

Topology

Figure 17. LDP Graceful Restart



Configuration

PE1

| | |
|--|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter Interface mode for loopback. |
| (config-if)# ip address 1.1.1.1/32 secondary | Assign secondary IP address. |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)# router ldp | Enter the Router LDP mode |
| (config-router)# router-id 1.1.1.1 | Assign router id for LDP |
| (config-router)# graceful-restart full | Enable graceful restart in LDP configure mode |
| (config-router)# graceful-restart timers neighbor-liveness 130 | Configuring graceful neighbor-liveness timer to 130 sec. |
| (config-router)# graceful-restart timers max-recovery 150 | Configuring graceful max recovery timer to 150 sec |
| (config-router)# exit | Exit Router LDP mode and return to Configure mode. |
| (config)#interface cd27/1 | Enter Interface mode for ce49. |
| (config-if)# ip address 20.20.20.1/24 | Assign IP address on ce49 interface. |

| | |
|---|---|
| (config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)# label-switching | Enable Label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)# router ospf | Enter the Router OSPF mode |
| (config-router)# network 1.1.1.1/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)# network 20.20.20.0/24 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)# redistribute connected | Configure redistribution connected. |
| (config-router)#exit | Exit from Router OSPF mode and enter into config mode |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#exit | Exit from configuration mode |

P

| | |
|---|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter Interface mode for loopback. |
| (config-if)# ip address 2.2.2.2/32 secondary | Assign secondary IP address. |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)# router ldp | Enter the Router LDP mode |
| (config-router)# router-id 2.2.2.2 | Assign router id for LDP |
| (config-router)# graceful-restart helper-only | Enable graceful restart helper in LDP configure mode |
| (config-router)# exit | Exit Router LDP mode and return to Configure mode. |
| (config)#interface xe12 | Enter Interface mode for ce0. |
| (config-if)# ip address 20.20.20.2/24 | Assign IP address on ce0 interface. |
| (config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)# label-switching | Enable Label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface xe22 | Enter Interface mode for xe13. |
| (config-if)# ip address 30.30.30.1/24 | Assign IP address on xe13 interface. |
| (config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)# label-switching | Enable Label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)# router ospf | Enter the Router OSPF mode |
| (config-router)# network 2.2.2.2/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)# network 20.20.20.0/24 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)# network 30.30.30.0/24 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)#exit | Exit from Router OSPF mode and enter into config |

| | |
|-----------------|---|
| | mode |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#exit | Exit from configuration mode |

PE2

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter Interface mode for loopback. |
| (config-if)# ip address 3.3.3.3/32 secondary | Assign secondary IP address. |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)# router ldp | Enter the Router LDP mode |
| (config-router)# router-id 3.3.3.3 | Assign router id for LDP |
| (config-router)# graceful-restart helper-only | Enable graceful restart helper in LDP configure mode |
| (config-router)# exit | Exit Router LDP mode and return to Configure mode. |
| (config)#interface xe5 | Enter Interface mode for xe13. |
| (config-if)# ip address 30.30.30.2/24 | Assign IP address on xe13 interface. |
| (config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)# label-switching | Enable Label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)# router ospf | Enter the Router OSPF mode |
| (config-router)# network 3.3.3.3/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)# network 30.30.30.0/24 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)# redistribute connected | Configure redistribution connected. |
| (config-router)#exit | Exit from Router OSPF mode and enter into config mode |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#exit | Exit from configuration mode |

Validation**PE1 Before LDP Graceful Restart**

```
PE1#show ip ospf neighbor
```

```
Total number of full neighbors: 1
```

```
OSPF process 0 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|------------|-----------|-------------|
| 2.2.2.2 | 1 | Full/DR | 00:00:33 | 20.20.20.2 | cd27/1 | 0 |

```
PE1#show running-config ldp
```

```
!
```

```
router ldp
```

```
router-id 1.1.1.1
```

```

graceful-restart full
graceful-restart timers neighbor-liveness 130
graceful-restart timers max-recovery 150
!
!
interface cd27/1
  enable-ldp ipv4
!
PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | cd27/1 | Passive | OPERATIONAL | 30 | 00:01:56 |

```

PE1#show mpls ldp parameter
Router ID           : 1.1.1.1
LDP Version         : 1
Fast-reroute Per-prefix : Disabled
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode  : Liberal
Label Control Mode    : Independent
Instance Loop Detection : Off
Request Retry         : Off
Propagate Release     : Disabled
Graceful Restart      : Enabled
Hello Interval        : 5
Targeted Hello Interval : 15
Hold time             : 15
Targeted Hold time    : 45
Keepalive Interval    : 10
Keepalive Timeout     : 30
Request retry Timeout : 5
Neighbor liveness Time : 130
Max recovery Time     : 150
Auto Targeted Hello Receipt : Disabled
Transport Address data :
  LabelSpace 0       : 1.1.1.1 (in use)
Import BGP routes    : No
Entropy-Label Capability : Disabled
Prefer-Tunnel-In-Tunnel : Disabled

```

```

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | Nexthop | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|---------------|---------|----------|----------|------------|-----|-----------|------|
| Intf | ELC | | Algo-Num | UpTime | | | | |
| L> | 2.2.2.2/32 | 1 | 2 | - | - | - | - | - |
| | - | N/A | 00:02:32 | | | | | |
| | | | 1 | - | | | | |
| | Yes | 3 | cd27/1 | No | 20.20.20.2 | - | - | - |
| L> | 3.3.3.3/32 | 2 | 4 | - | - | - | - | - |
| | - | N/A | 00:02:32 | | | | | |
| | | | 3 | - | | | | |
| | Yes | 24320 | cd27/1 | No | 20.20.20.2 | - | - | - |
| L> | 30.30.30.0/24 | 3 | 2 | - | - | - | - | - |
| | - | N/A | 00:02:32 | | | | | |
| | | | 1 | - | | | | |
| | Yes | 3 | cd27/1 | No | 20.20.20.2 | - | - | - |


```
PE1#show ldp session multicast count
```

```
-----
session up count   : 1
session down count : 0
total count        : 1
-----
```

```
PE1#show ldp session all count
```

```
-----
session up count   : 1
session down count : 0
total count        : 1
-----
```

```
PE1#show ldp adjacency count
```

```
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received
```

| Type | Entity | Adjacency |
|-----------|--------|-----------|
| Multicast | 1 | 1 |
| Targeted | 0 | 0 |
| Total | 1 | 1 |

```
-----
```

```
PE1#show mpls forwarding-table detail
```

```
FEC prefix: 2.2.2.2/32, FTN-ID: 2
```

```
Owner: LDP, FTN type: REGULAR, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
ECMP nhlfe-ix: 4
Primary : Out-Label: 3, Out-Intf: cd27/1, Nexthop: 20.20.20.2
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No
```

```
FEC prefix: 3.3.3.3/32, FTN-ID: 1
```

```
Owner: LDP, FTN type: REGULAR, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
ECMP nhlfe-ix: 3
Primary : Out-Label: 24321, Out-Intf: cd27/1, Nexthop: 20.20.20.2
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No
```

```
FEC prefix: 30.30.30.0/24, FTN-ID: 3
```

```
Owner: LDP, FTN type: REGULAR, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
ECMP nhlfe-ix: 4
Primary : Out-Label: 3, Out-Intf: cd27/1, Nexthop: 20.20.20.2
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No
```

PE1 after Graceful Restart

```
PE1#wr
```

```
Building configuration...
```

```
[OK]
```

```
PE1#restart ldp graceful
```

```
% Warning : LDP process will stop and needs to restart manually,
You may loose ldp configuration,if not saved
```

```

Proceed for graceful restart? (y/n):y
Managed module 'ldp' is busy or down
PE1#2025 Jul 03 07:42:21.922 : PE1 : CML : CRITI : Module ldp
disc                                     onnected with CML
2025 Jul 03 07:42:21.922 : PE1 : LDP : CRITI : [LDP_SESSION_DOWN_
2]                                     : Clearing up session on
interface cd27/1 with peer 2.2.2.2, last
do                                     wn reason None

```

Verify the sessions and running configs for LDP both should be Empty

```
PE1#show ldp session
```

```
PE1#show mpls ldp parameter
```

```
PE1#show mpls forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
        B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
        (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|---------------|---------|----------|-----------|-----|------------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> p | 2.2.2.2/32 | 1 | 2 | - | - | - | - |
| | - | N/A | 00:05:47 | | | | |
| | Yes | 3 | cd27/1 | No | p | 20.20.20.2 | - |
| L> p | 3.3.3.3/32 | 2 | 4 | - | - | - | - |
| | - | N/A | 00:05:47 | | | | |
| | Yes | 24320 | cd27/1 | No | p | 20.20.20.2 | - |
| L> p | 30.30.30.0/24 | 3 | 2 | - | - | - | - |
| | - | N/A | 00:05:47 | | | | |
| | Yes | 3 | cd27/1 | No | p | 20.20.20.2 | - |

```
PE1#show mpls forwarding-table detail
```

```
FEC prefix: 2.2.2.2/32, FTN-ID: 2
```

```

Owner: LDP, FTN type: REGULAR, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
ECMP nhlfe-ix: 4
Primary : Out-Label: 3, Out-Intf: cd27/1, Nexthop: 20.20.20.2
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

```

```
FEC prefix: 3.3.3.3/32, FTN-ID: 1
```

```

Owner: LDP, FTN type: REGULAR, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
ECMP nhlfe-ix: 3
Primary : Out-Label: 24321, Out-Intf: cd27/1, Nexthop: 20.20.20.2
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

```

```
FEC prefix: 30.30.30.0/24, FTN-ID: 3
```

```

Owner: LDP, FTN type: REGULAR, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
ECMP nhlfe-ix: 4
Primary : Out-Label: 3, Out-Intf: cd27/1, Nexthop: 20.20.20.2
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

```

```
PE1#show nsm ldp forwarding-timer
```

```

Protocol-Name  GR-State  Time Remaining (sec)  Disconnected-time
LDP           ACTIVE    119                  2025/07/03 07:47:02 PE1

```

```

PE1#ping 3.3.3.3
Press CTRL+C to exit
PING 3.3.3.3 (3.3.3.3) 100(128) bytes of data.
108 bytes from 3.3.3.3: icmp_seq=1 ttl=64 time=0.687 ms
108 bytes from 3.3.3.3: icmp_seq=2 ttl=64 time=0.585 ms
108 bytes from 3.3.3.3: icmp_seq=3 ttl=64 time=1.18 ms
108 bytes from 3.3.3.3: icmp_seq=4 ttl=64 time=0.682 ms
108 bytes from 3.3.3.3: icmp_seq=5 ttl=64 time=1.18 ms
108 bytes from 3.3.3.3: icmp_seq=6 ttl=64 time=1.23 ms
108 bytes from 3.3.3.3: icmp_seq=7 ttl=64 time=0.879 ms
108 bytes from 3.3.3.3: icmp_seq=8 ttl=64 time=0.798 ms
108 bytes from 3.3.3.3: icmp_seq=9 ttl=64 time=0.937 ms

--- 3.3.3.3 ping statistics ---
9 packets transmitted, 9 received, 0% packet loss, time 8165ms
rtt min/avg/max/mdev = 0.585/0.906/1.231/0.228 ms

```

PE1 Restarting the LDP process

```

PE1#start-shell
bash-5.2$ su -
Password:
root@PE1:~# cd /usr/local/sbin/
root@PE1:/usr/local/sbin# ./ldpd -d
root@PE1:/usr/local/sbin# exit
logout
bash-5.2$ exit
exit
PE1#show running-config ldp
!
router ldp
 router-id 1.1.1.1
 graceful-restart full
 graceful-restart timers neighbor-liveness 130
 graceful-restart timers max-recovery 150
!
!
interface cd27/1
 enable-ldp ipv4
!
PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|--------|
| | 2.2.2.2 | cd27/1 | Passive | OPERATIONAL | 30 | |

```

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

| FTN-ECMP LDP: Disabled, SR: Disabled | | | | | | | | | |
|--------------------------------------|------------|---------|----------|----------|------------|-----|-----------|------|---|
| Code | FEC | Nexthop | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- | |
| Intf | ELC | | Algo-Num | UpTime | | | | | |
| L> | 2.2.2.2/32 | 1 | 5 | - | - | - | - | - | - |
| | - | N/A | 00:00:03 | | | | | | |
| | | | 1 | - | | | | | |
| | Yes | 3 | cd27/1 | No | 20.20.20.2 | | - | - | |

```

L> 3.3.3.3/32      2      10      -      -      -      -      -
    -      N/A      00:00:03      6      -
    Yes 24321      cd27/1      No      20.20.20.2      -      -
L> 30.30.30.0/24   3      5      -      -      -      -      -
    -      N/A      00:00:03      1      -
    Yes 3      cd27/1      No      20.20.20.2      -      -

```

PE1#show mpls ldp parameter

```

Router ID      : 1.1.1.1
LDP Version    : 1
Fast-reroute Per-prefix : Disabled
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode : Liberal
Label Control Mode : Independent
Instance Loop Detection : Off
Request Retry   : Off
Propagate Release : Disabled
Graceful Restart : Enabled
Hello Interval  : 5
Targeted Hello Interval : 15
Hold time       : 15
Targeted Hold time : 45
Keepalive Interval : 10
Keepalive Timeout : 30
Request retry Timeout : 5
Neighbor liveness Time : 130
Max recovery Time : 150
Auto Targeted Hello Receipt : Disabled
Transport Address data :
  Labelspace 0 : 1.1.1.1 (in use)
Import BGP routes : No
Entropy-Label Capability : Disabled
Prefer-Tunnel-In-Tunnel : Disabled

```

P Before Restart LDP

P#show ip ospf neighbor

Total number of full neighbors: 2

OSPF process 0 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 1.1.1.1 | 1 | Full/Backup | 00:00:38 | 20.20.20.1 | xe12 | 0 |
| 3.3.3.3 | 1 | Full/DR | 00:00:37 | 30.30.30.2 | xe22 | 0 |

P#show running-config ldp

```

!
router ldp
  router-id 2.2.2.2
  graceful-restart helper-only
!
!
interface xe12
  enable-ldp ipv4
!
interface xe22
  enable-ldp ipv4
!

```

P#show mpls ldp parameter

```

Router ID      : 2.2.2.2

```

```

LDP Version           : 1
Fast-reroute Per-prefix : Disabled
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode    : Liberal
Label Control Mode      : Independent
Instance Loop Detection : Off
Request Retry           : Off
Propagate Release       : Disabled
Graceful Restart        : Enabled
GracefulRestartHelpermode: Enabled
Hello Interval          : 5
Targeted Hello Interval : 15
Hold time               : 15
Targeted Hold time      : 45
Keepalive Interval      : 10
Keepalive Timeout       : 30
Request retry Timeout    : 5
Neighbor liveness Time  : 120
Max recovery Time       : 120
Auto Targeted Hello Receipt : Disabled
Transport Address data  :
  Labelspace 0         : 2.2.2.2 (in use)
Import BGP routes       : No
Entropy-Label Capability : Disabled
Prefer-Tunnel-In-Tunnel : Disabled
P#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code  FEC          FTN-ID  Nhlfe-ID  Tunnel-ID  Pri  Out-Label  Out-
Intf  ELC          Nexthop  Algo-Num  UpTime
L>    1.1.1.1/32    2        5         -         -    -         -
      -            N/A      00:04:08
      4           -
      Yes 3        xe12     No        20.20.20.1  -    -         -
L>    3.3.3.3/32    1        11        -         -    -         -
      -            N/A      00:21:13
      10         -
      Yes 3        xe22     No        30.30.30.2  -    -         -

P#show ldp session multicast count
-----
session up count   : 2
session down count : 0
total count        : 2
-----

P#show ldp session all count
-----
session up count   : 2
session down count : 0
total count        : 2
-----

P#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity      Adjacency
-----

```

```

Multicast  2          2
Targeted   0          0
Total      2          2
-----

```

```

P#show mpls forwarding-table detail
FEC prefix: 1.1.1.1/32, FTN-ID: 2
  Owner: LDP, FTN type: REGULAR, State: Installed
  Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
  LSP-ID: N/A, LSP-type: Primary
  ECMP nhlfex: 5
  Primary : Out-Label: 3, Out-Intf: xe12, Nexthop: 20.20.20.1
  Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
  ELC: No

```

```

FEC prefix: 3.3.3.3/32, FTN-ID: 1
  Owner: LDP, FTN type: REGULAR, State: Installed
  Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
  LSP-ID: N/A, LSP-type: Primary
  ECMP nhlfex: 2
  Primary : Out-Label: 3, Out-Intf: xe22, Nexthop: 30.30.30.2
  Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0

```

P After LDP Graceful Restart

```

P#wr
Building configuration...
[OK]
P#restart ldp graceful
% Warning : LDP process will stop and needs to restart manually,
You may loose ldp configuration,if not saved
Proceed for graceful restart? (y/n):y
Managed module 'ldp' is busy or down
P#2025 Jun 04 11:56:09.525 : P : CML : CRITI : Module ldp
di                                     sconnected with CML
2025 Jun 04 11:56:09.525 : P : LDP : CRITI : [LDP_SESSION_DOWN_
2                                     ]: Clearing up session on
interface xe12 with peer 1.1.1.1, last                                     n reason None
dow
2025 Jun 04 11:56:09.525 : P : LDP : CRITI : [LDP_SESSION_DOWN_
2                                     ]: Clearing up session on
interface xe22 with peer 3.3.3.3, last                                     n reason None
dow

```

Verify the sessions and running configs for LDP both should be Empty

```

P#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 0 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface      Instance ID
1.1.1.1        1     Full/Backup     00:00:32   20.20.20.1    xe12           0
3.3.3.3        1     Full/DR         00:00:31   30.30.30.2    xe22           0

```

```

P#show run ldp

```

```

!
!

```

```

P#show ldp session

```

```

P#show mpls forwarding-table detail
FEC prefix: 1.1.1.1/32, FTN-ID: 2
  Owner: LDP, FTN type: REGULAR, State: Installed
  Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
  LSP-ID: N/A, LSP-type: Primary

```

```

ECMP nhlfe-ix: 5
Primary : Out-Label: 3, Out-Intf: xe12, Nexthop: 20.20.20.1
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

P#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code   FEC          FTN-ID   Nhife-ID  Tunnel-ID   Pri   Out-Label   Out-
Intf   ELC          Nexthop   Algo-Num  UpTime
L> p 1.1.1.1/32      2        5         -           -     -           -
      -           N/A      00:04:39
      4
      Yes 3        xe12      No p      20.20.20.1   -     -

```

PE2 Before LDP Graceful Restart

```

PE2#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 0 VRF(default):
Neighbor ID   Pri   State           Dead Time   Address        Interface      Instance ID
2.2.2.2       1    Full/Backup     00:00:34   30.30.30.1    xe5            0

PE2#show running-config ldp
!
router ldp
  router-id 3.3.3.3
  graceful-restart helper-only
!
!
interface xe5
  enable-ldp ipv4
!

PE2#show mpls ldp parameter
Router ID           : 3.3.3.3
LDP Version         : 1
Fast-reroute Per-prefix : Disabled
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode  : Liberal
Label Control Mode    : Independent
Instance Loop Detection : Off
Request Retry         : Off
Propagate Release     : Disabled
Graceful Restart      : Enabled
GracefulRestartHelpermode: Enabled
Hello Interval        : 5
Targeted Hello Interval : 15
Hold time             : 15
Targeted Hold time    : 45
Keepalive Interval     : 10
Keepalive Timeout      : 30
Request retry Timeout  : 5
Neighbor liveness Time : 120
Max recovery Time      : 120
Auto Targeted Hello Receipt : Disabled
Transport Address data :
  Labelspace 0        : 3.3.3.3 (in use)
Import BGP routes     : No

```

```
Entropy-Label Capability : Disabled
Prefer-Tunnel-In-Tunnel : Disabled
```

```
PE2#show mpls forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP
```

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|---------------|---------|----------|-----------|-----|------------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 1 | 8 | - | - | - | - |
| | - | N/A | 00:21:41 | | | | |
| | | | 7 | - | | | |
| | Yes | 24321 | xe5 | No | | 30.30.30.1 | - |
| L> | 2.2.2.2/32 | 2 | 9 | - | - | - | - |
| | - | N/A | 00:21:41 | | | | |
| | | | 3 | - | | | |
| | Yes | 3 | xe5 | No | | 30.30.30.1 | - |
| L> | 20.20.20.0/24 | 3 | 9 | - | - | - | - |
| | - | N/A | 00:21:41 | | | | |
| | | | 3 | - | | | |
| | Yes | 3 | xe5 | No | | 30.30.30.1 | - |

```
PE2#show ldp session multicast count
```

```
-----
session up count   : 1
session down count : 0
total count        : 1
-----
```

```
PE2#show ldp session all count
```

```
-----
session up count   : 1
session down count : 0
total count        : 1
-----
```

```
PE2#show ldp adjacency count
```

```
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received
```

| Type | Entity | Adjacency |
|-----------|--------|-----------|
| Multicast | 1 | 1 |
| Targeted | 0 | 0 |
| Total | 1 | 1 |

```
PE2#show mpls forwarding-table detail
```

```
FEC prefix: 1.1.1.1/32, FTN-ID: 1
```

```
Owner: LDP, FTN type: REGULAR, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
ECMP nhlfe-ix: 2
Primary : Out-Label: 24320, Out-Intf: xe5, Nexthop: 30.30.30.1
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No
```

```
FEC prefix: 2.2.2.2/32, FTN-ID: 2
```

```
Owner: LDP, FTN type: REGULAR, State: Installed
```



```

Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
ECMP nhlfex: 4
Primary : Out-Label: 3, Out-Intf: xe5, Nexthop: 30.30.30.1
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

FEC prefix: 20.20.20.0/24, FTN-ID: 3
Owner: LDP, FTN type: REGULAR, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
ECMP nhlfex: 4
Primary : Out-Label: 3, Out-Intf: xe5, Nexthop: 30.30.30.1
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

```

PE2 After LDP Graceful Restart

```

PE2#wr
Building configuration...
[OK]
PE2#restart ldp graceful
% Warning : LDP process will stop and needs to restart manually,
You may loose ldp configuration,if not saved
Proceed for graceful restart? (y/n):y
Graceful restart cannot be performed when helper mode is active
PE2#2025 May 31 14:37:43.932 : PE2 : LDP : CRITI : [LDP_SESSION_
                                DOWN_2]: Clearing up session
on interface xe5 with peer 2.2.2.2,
las                                t down reason Session Read
Error
2025 May 31 14:37:58.693 : PE2 : LDP : CRITI : [LDP_SESSION_DOWN_2]: Clearing up session on interface
xe5 with peer 2.2.2.2, session down reason Session Read Error

PE2#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 0 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface       Instance ID
2.2.2.2        1     Full/Backup     00:00:36   30.30.30.1    xe5             0

Verify the sessions and running configs for LDP both should be Empty

PE2#show running-config ldp
!
!

PE2#show ldp session

```

PE2 Restarting the LDP process

```

PE2#start-shell
bash-5.2$ su -
Password:
root@PE2:~# cd /usr/local/sbin/
root@PE2:/usr/local/sbin# ./ldpd -d
Other ldpd may be running. Please check it
root@PE2:/usr/local/sbin# exit
logout
bash-5.2$
bash-5.2$ exit
exit

```

```
PE2#sh mpls forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
        B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
        (m) - FTN mapped over multipath transport, (e) - FTN is ECMP
```

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | Nexthop | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|---------------|---------|----------|----------|------------|-----|-----------|------|
| Intf | ELC | | Algo-Num | UpTime | | | | |
| L> | 1.1.1.1/32 | 1 | 2 | - | - | - | - | - |
| | - | N/A | 00:01:05 | - | - | - | - | - |
| | Yes | 24320 | xe5 | No | 30.30.30.1 | - | - | - |
| L> | 2.2.2.2/32 | 2 | 4 | - | - | - | - | - |
| | - | N/A | 00:01:05 | - | - | - | - | - |
| | Yes | 3 | xe5 | No | 30.30.30.1 | - | - | - |
| L> | 20.20.20.0/24 | 3 | 4 | p | - | - | - | - |
| | - | N/A | 00:01:05 | - | - | - | - | - |
| | Yes | 3 | xe5 | No | 30.30.30.1 | - | - | - |

```
PE2#sh run ldp
```

```
!
router ldp
router-id 3.3.3.3
graceful-restart helper-only
!
!
interface xe5
enable-ldp ipv4
!
```

LDP Tunneling over RSVP-TE

Overview

LDP-over-RSVP-TE tunneling is a technique used in MPLS networks to combine the strengths of Label Distribution Protocol (LDP) and Resource Reservation Protocol Traffic Engineering (RSVP-TE). This approach allows LDP Label Switched Paths (LSPs) to be encapsulated within RSVP-TE LSPs, providing enhanced traffic engineering capabilities while maintaining operational simplicity.

Feature Characteristics

LDP-over-RSVP-TE facilitates the integration of LDP LSPs within RSVP-TE tunnels, leveraging the strengths of both protocols. It harnesses RSVP-TE's traffic engineering capabilities for path computation, bandwidth reservation, and quality of service (QoS) provisioning. Ingress nodes execute FEC resolution to designate the suitable RSVP-TE tunnel for tunneling LDP LSPs, establishing hierarchical LSPs with RSVP-TE as the outer label and LDP as the inner label.

Benefits

LDP-over-RSVP-TE offers significant benefits are:

- **Advanced Traffic Engineering:** By leveraging RSVP-TE's advanced traffic engineering mechanisms, LDP-over-RSVP-TE enables efficient path computation, bandwidth reservation, and Quality of Service (QoS) provisioning.
- **Simplified Network Topology:** eliminates the need for a full mesh of intra-area RSVP LSPs (Label Switched Paths) between PE (Provider Edge) nodes.
- **Enhanced Resilience with Fast Reroute (FRR):** Inherit RSVP-TE's Fast Reroute (FRR) capabilities. This means that in case of link or node failures, the network can quickly reroute traffic along pre-established backup paths
- **Flexible Hierarchical LSP Design:** Provides flexibility in network design by allowing for hierarchical LSPs (Label Switched Paths) where RSVP-TE serves as the outer label and LDP as the inner label.

Prerequisites

Before configuring this feature, ensure the following:

- A functional MPLS network with support for both LDP and RSVP-TE protocols.
- Network devices (routers or switches) capable of supporting LDP and RSVP-TE functionalities.

Limitations

The limitations are:

- Tunneling over inter-domain IGP area is not supported.
- LDP LSP tunneling over RSVP multipath is not supported.
- MPLS trace route is not supported in LDP-over-RSVP tunneling path.
- Dynamic TLDP sessions are not supported, TLDP session has to be explicitly configured.

- LFA and/or RLFA protection is not supported for LDP-over-RSVP tunnels.
- LDP over RSVP is supported only in a linear topology, where LDP traffic traversing the RSVP domain cannot be looped back toward the source LDP domain due to the presence of an IGP path that would cause such a loop.
- In QMX devices being LSR node, Entropy Label (EL) and Entropy Label Indicator (ELI) cannot be inserted after RSVP label for "LDP over RSVP" use case.
- LDP over RSVP (LoR) FTN updates are not hitless if the RSVP trunk undergoes an MBB sequence.
- LDP over RSVP (LoR) is supported only for OSPF intra-area routes. The Type 3 and Type 5 routes are not considered for LoR.


Configuration for LDP Tunneling Over RSVP

Configure various nodes within the topology to set up a LDP Tunneling over RSVP session.

Topology

This sample topology provides basic connectivity and routing between the devices.

Figure 18. LDP Tunneling over RSVP Configuration



Configure LDP Tunneling over RSVP

PE1 Router

Follow the steps to configure the LDP tunneling over RSVP on PE1 router:

1. Configure the loopback interface with an IP address.

```
PE1(config)#interface lo
PE1(config-if)#ip address 10.12.96.1/32 secondary
```

2. Configure the global LDP parameters including the router ID and transport address.

```
PE1(config)#router ldp
PE1(config-router)#router-id 10.12.96.1
PE1(config-router)#transport-address ipv4 10.12.96.1
```

3. Configure global RSVP parameters.

```
PE1(config)#router rsvp
```

4. Configure the interface facing the network side with an IP address, enable label switching, and enable LDP.

```
PE1(config)#interface xe1
PE1(config-if)#ip address 10.1.1.1/24
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
```

5. If using ISIS as the Interior Gateway Protocol (IGP), configure ISIS parameters including traffic engineering.

```
PE1(config)#router isis ISIS-IGP
PE1(config-router)#is-type level-1
PE1(config-router)#metric-style wide
PE1(config-router)#mpls traffic-eng router-id 10.12.96.1
PE1(config-router)#mpls traffic-eng level-1
PE1(config-router)#capability cspf
PE1(config-router)#dynamic-hostname
PE1(config-router)#net 49.0000.0000.0001.00
PE1(config-router)#exit

PE1(config)#interface lo
PE1(config-if)#ip router isis ISIS-IGP
PE1(config)#interface xe1
PE1(config-if)# ip router isis ISIS-IGP
```

P1 Router

Follow the steps to configure the LDP tunneling over RSVP on P1 router:

1. Configure the loopback interface with an IP address.

```
P1(config)#interface lo
P1(config-if)# ip address 10.12.96.2/32 secondary
```

2. Configure the global TLDP parameters including the router ID and transport address.

```
P1(config)#router ldp
P1(config-router)#router-id 10.12.96.2
P1(config-router)# targeted-peer ipv4 10.12.96.4
P1(config-router-targeted-peer)# tunneling
P1(config-router-targeted-peer)#exit
P1(config-router)# transport-address ipv4 10.12.96.2
```

3. Configure LDP to prefer tunneling over RSVP.

```
P1(config)#router ldp
P1(config-router)# prefer-tunnel-in-tunnel rsvp
```

4. Configure global RSVP parameters.

```
P1(config)#router rsvp
```

5. Configure a RSVP trunk towards the neighbor router (assuming 10.12.96.4 is the neighbor) and enable ldp-tunneling to allow tunneling LDP LSPs.

```
P1(config)# rsvp-trunk t1 ipv4
P1(config-trunk)#to 10.12.96.4
P1(config-trunk)#ldp-tunneling
```

6. Configure the interface facing the network side with an IP address, enable label switching, and enable LDP and RSVP.

- For interface xe1:

```
P1(config)#interface xe1
P1(config-if)#ip address 10.1.1.2/24
P1(config-if)#label-switching
P1(config-if)#enable-ldp ipv4
```

- For interface xe2:

```
P1(config)#interface xe2
```

```
P1(config-if)#ip address 20.1.1.1/24
P1(config-if)#label-switching
P1(config-if)#enable-rsvp
```

7. If using ISIS as the Interior Gateway Protocol (IGP), configure ISIS parameters including traffic engineering.

```
P1(config)#router isis ISIS-IGP
P1(config-router)#is-type level-1
P1(config-router)#metric-style wide
P1(config-router)#mpls traffic-eng router-id 10.12.96.2
P1(config-router)#mpls traffic-eng level-1
P1(config-router)#capability cspf
P1(config-router)#dynamic-hostname
P1(config-router)#net 49.0000.0000.0002.00
P1(config-router)#exit

P1(config)#interface lo
P1(config-if)# ip router isis ISIS-IGP
P1(config)#interface xe1
P1(config-if)# ip router isis ISIS-IGP
P1(config)#interface xe2
P1(config-if)# ip router isis ISIS-IGP
```

P2 Router

Follow the steps to configure the LDP tunneling over RSVP on P2 router:

1. Configure the loopback interface with an IP address.

```
P2(config)#interface lo
P2(config-if)# ip address 10.12.96.3/32 secondary
```

2. Configure the global LDP parameters including the router ID and transport address.

```
P2(config)#router ldp
P2(config-router)#router-id 10.12.96.3
P2(config-router)# transport-address ipv4 10.12.96.3
```

3. Configure global RSVP parameters.

```
P2(config)#router rsvp
```

4. Configure the interface facing the network side with an IP address, enable label switching, and enable RSVP.

- For interface xe2:

```
P2(config)#interface xe2
P2(config-if)#ip address 20.1.1.2/24
P2(config-if)#label-switching
P2(config-if)#enable-rsvp
```

- For interface xe3:

```
P2(config)#interface xe3
P2(config-if)#ip address 30.1.1.1/24
P2(config-if)#label-switching
P2(config-if)#enable-rsvp
```

5. If using ISIS as the Interior Gateway Protocol (IGP), configure ISIS parameters including traffic engineering.

```
P2(config)#router isis ISIS-IGP
P2(config-router)#is-type level-1
P2(config-router)#metric-style wide
P2(config-router)#mpls traffic-eng router-id 10.12.96.3
P2(config-router)#mpls traffic-eng level-1
```

```
P2(config-router)#capability cspf
P2(config-router)#dynamic-hostname
P2(config-router)#net 49.0000.0000.0003.00
P2(config-router)#exit

P2(config)#interface lo
P2(config-if)# ip router isis ISIS-IGP
P2(config)#interface xe2
P2(config-if)# ip router isis ISIS-IGP
P2(config)#interface xe3
P2(config-if)# ip router isis ISIS-IGP
```

P3 Router

Follow the steps to configure the LDP tunneling over RSVP on P3 router:

1. Configure the loopback interface with an IP address.

```
P3(config)#interface lo
P3(config-if)# ip address 10.12.96.4/32 secondary
```

2. Configure the global LDP parameters including the router ID and transport address.

```
P3(config)#router ldp
P3(config-router)#router-id 10.12.96.4
P3(config-router)# targeted-peer ipv4 10.12.96.2
P3(config-router-targeted-peer)# tunneling
P3(config-router-targeted-peer)#exit
P3(config-router)# transport-address ipv4 10.12.96.4
```

3. Configure global RSVP parameters.

```
P3(config)#router rsvp
```

4. Configure prefix lists.

```
P3(config)# ip prefix-list fec_list
P3(config-ip-prefix-list)# seq 5 permit 10.12.96.5/32
P3(config)# ip prefix-list peer_list
P3(config-ip-prefix-list)# seq 5 permit 10.12.96.2/32
```

5. Configure prefix lists to control label advertisement between peers.

```
P3(config)# router ldp
P3(config-router)# advertise-labels for fec_list to peer_list
```

6. Configure the interface facing the network side with an IP address, enable label switching, and enable RSVP.

- For interface xe3:

```
P3(config)#interface xe3
P3(config-if)#ip address 30.1.1.2/24
P3(config-if)#label-switching
P3(config-if)#enable-rsvp
```

- For interface xe4:

```
P3(config)#interface xe4
P3(config-if)#ip address 40.1.1.1/24
P3(config-if)#label-switching
P3(config-if)#enable-ldp ipv4
```

7. If using ISIS as the Interior Gateway Protocol (IGP), configure ISIS parameters including traffic engineering.

```
P3(config)#router isis ISIS-IGP
P3(config-router)#is-type level-1
P3(config-router)#metric-style wide
P3(config-router)#mpls traffic-eng router-id 10.12.96.4
P3(config-router)#mpls traffic-eng level-1
P3(config-router)#capability cspf
P3(config-router)#dynamic-hostname
P3(config-router)#net 49.0000.0000.0004.00
P3(config-router)#exit

P3(config)#interface lo
P3(config-if)# ip router isis ISIS-IGP
P3(config)#interface xe3
P3(config-if)# ip router isis ISIS-IGP
P3(config)#interface xe4
P3(config-if)# ip router isis ISIS-IGP
```

PE2 Router

Follow the steps to configure the LDP tunneling over RSVP on PE2 router:

1. configure the loopback interface with an IP address.

```
PE2(config)#interface lo
PE2(config-if)# ip address 10.12.96.5/32 secondary
```

2. Configure the global LDP parameters including the router ID and transport address.

```
PE2(config)#router ldp
PE2(config-router)#router-id 10.12.96.5
PE2(config-router)# transport-address ipv4 10.12.96.5
```

3. Configure the interface facing the network side with an IP address, enable label switching, and enable LDP.

```
PE2(config)#interface xe4
PE2(config-if)#ip address 40.1.1.2/24
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
```

4. If using ISIS as the Interior Gateway Protocol (IGP), configure ISIS parameters including traffic engineering.

```
PE2(config)#router isis ISIS-IGP
PE2(config-router)#is-type level-1
PE2(config-router)#metric-style wide
PE2(config-router)#mpls traffic-eng router-id 10.12.96.5
PE2(config-router)#mpls traffic-eng level-1
PE2(config-router)#capability cspf
PE2(config-router)#dynamic-hostname
PE2(config-router)#net 49.0000.0000.0005.00
PE2(config-router)#exit

PE2(config)#interface lo
PE2(config-if)# ip router isis ISIS-IGP
PE2(config)#interface xe4
PE2(config-if)# ip router isis ISIS-IGP
```

Running Configuration

P1 Router

Follow the steps to configure the LDP tunneling over RSVP on P1 router using snippet:


```
P1#show running-config isis
!
router isis ISIS-IGP
 is-type level-1
 metric-style wide
 mpls traffic-eng router-id 10.12.96.2
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 net 49.0000.0000.0002.00
!
interface xe1
 ip router isis ISIS-IGP
!
interface lo
 ip router isis ISIS-IGP
!
interface xe2
 ip router isis ISIS-IGP
!
P1#show running-config ldp
!
router ldp
 router-id 10.12.96.2
 prefer-tunnel-in-tunnel rsvp
 targeted-peer ipv4 10.12.96.4
 tunneling
  exit-targeted-peer-mode
 transport-address ipv4 10.12.96.2
!
!
interface xe1
 enable-ldp ipv4
!
P1#show running-config rsvp
!
router rsvp
!
!
interface xe2
 enable-rsvp
!
!
!
!
!
rsvp-trunk t1 ipv4
 to 10.12.96.4
 ldp-tunneling
!
```

P3 Router

Follow the steps to configure the LDP tunneling over RSVP on P3 router using snippet:

```
P3#show running-config ldp
!
router ldp
 router-id 10.12.96.4
 targeted-peer ipv4 10.12.96.2
 tunneling
  exit-targeted-peer-mode
 transport-address ipv4 10.12.96.4
 advertise-labels for fec_list to peer_list
!
!
interface xe4
```

```

enable-ldp ipv4
!
P3#

P3#show running-config prefix-list
!
ip prefix-list fec_list
seq 5 permit 10.12.96.5/32
!
ip prefix-list peer_list
seq 5 permit 10.12.96.2/32
!

```

Validation

Validation on P1 Node

```

P1#show ldp
Router ID           : 10.12.96.2
LDP Version         : 1
Fast-reroute Per-prefix : Disabled
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode   : Liberal
Label Control Mode     : Independent
Instance Loop Detection : Off
Request Retry          : Off
Propagate Release      : Disabled
Graceful Restart       : Disabled
Hello Interval         : 5
Targeted Hello Interval : 15
Hold time              : 15
Targeted Hold time     : 45
Keepalive Interval     : 10
Keepalive Timeout      : 30
Request retry Timeout   : 5
Auto Targeted Hello Receipt : Disabled
Transport Address data :
  Labelspace 0        : 10.12.96.2 (in use)
Import BGP routes      : No
Entropy-Label Capability : Disabled
Prefer-Tunnel-In-Tunnel : Enabled

P1#show ldp adjacency count
Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type      Entity  Adjacency
-----
Multicast  1       1
Targeted   1       1
Total      2       2
-----

P1#show ldp session targeted count
-----
session up count   : 1
session down count : 0
total count        : 1
-----

P1#show ldp session all count
-----
session up count   : 2
session down count : 0
total count        : 2

```

Pl#show ldp session

Codes: m - MD5 password is not set/unset.
 g - GR configuration not set/unset.
 t - TCP MSS not set/unset.
 Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 10.12.96.4 | xe2 | Passive | OPERATIONAL | 30 | 00:46:52 |
| | 10.12.96.1 | xe1 | Active | OPERATIONAL | 30 | 00:46:55 |

Pl#show rsvp session detail

Ingress (Primary)

10.12.96.4

From: 10.12.96.2, LSPstate: Up, LSPname: t1-Primary
 Ingress FSM state: Operational
 Establishment Time: 0s 15ms
 Setup priority: 7, Hold priority: 0, HOP limit: 255
 CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
 LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: ISIS
 IGP-Shortcut: Disabled, LSP metric: 20
 LDP Tunneling : Enabled
 LSP Protection: None
 Label in: -, Label out: 24320,
 Tspec rate: 0, Fspec rate: 0
 Policer: Not Configured
 Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 10.12.96.2
 Bind value: 0, Oper state: NA, Alloc mode: NA
 Downstream: 20.1.1.2, xe2
 Path refresh: 30 seconds (RR enabled) (due in 27203 seconds)
 Resv lifetime: 157 seconds (due in 154 seconds)
 Retry count: 0, intrvl: 30 seconds
 RRO re-use as ERO: Disabled
 Label Recording: Disabled
 Admin Groups: none
 Configured Path: none
 Session Explicit Route Detail :
 20.1.1.2/32 strict
 30.1.1.2/32 strict
 Record route:

 IP Address Label

<self>

20.1.1.2

30.1.1.2

Style: Shared Explicit Filter

Traffic type: controlled-load

Minimum Path MTU: 1500

Recorded Time : N/A

Current Error:

Code : None, Value : None

Originated Node : None, Recorded Time : N/A

Last Signaled Error:

Code : None, Value : None

Originated Node : None, Recorded Time : N/A

Trunk Type: mpls

Pl#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled

```

Code   FEC           FTN-ID   Nhlfe-ID   Tunnel-ID   Pri   Out-Label   Out-
Intf   ELC           Nexthop   UpTime
L>    10.12.96.1/32    2        39         -           -     -           -
      -              00:31:26    38         -           -
      Yes    3        xe1        No        10.1.1.1    -
R
(t)>   10.12.96.4/32    1        9          5001        Yes   25600        xe2        No        2
2.1.1.1 00:01:19
L>    10.12.96.5/32    3        11         -           -     -           -
      -              00:01:19    10         -           Yes   26244        -
      No        10.12.96.4    -

```

(via rsvp tunnel-id 5001,

nhlfe_ix 9, label 25600)

Pl#

Pl#

Pl#

Pl#show ldp tunneling

```

Tunnel Name       : t1
Tunnel Endpoint   : 10.12.96.4/32
Tunnel Cost       : 20
Tunnel Owner      : RSVP
Tunnel Status     : Up

```

| FEC | Upstream-Peer | In-Label | Out-Label |
|---------------|---------------|----------|-----------|
| 10.12.96.5/32 | 10.12.96.4 | 24969 | 24966 |
| | 10.12.96.1 | 24963 | 24966 |

Total FEC tunneled by t1 : 1

Pl#show ldp tunneling-fec

| FEC | Tunnel-name | Tunnel-endpoint | Upstream-Peer | In-label | Out- |
|---------------|-------------|-----------------|---------------|----------|-------|
| 10.12.96.5/32 | t1 | 10.12.96.4/32 | 10.12.96.4 | 24969 | 24966 |
| | | | 10.12.96.1 | 24963 | 24966 |

Total LDP Tunneled FEC : 1

Pl#show mpls forwarding-table

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

FTN-ECMP LDP: Disabled, SR: Disabled

```

Code   FEC           FTN-ID   Nhlfe-ID   Tunnel-ID   Pri   Out-Label   Out-
Intf   ELC           Nexthop   Algo-Num   UpTime
L>    10.12.96.1/32    1        2          -           -     -           -
      -              N/A      00:57:59    1          -           -
      Yes    3        xe1        No        10.1.1.1    -
R
(t)>   10.12.96.4/32    2        4          5001        Yes   24320        xe2        No        2
0.1.1.2 N/A      00:57:35
L>    10.12.96.5/32    3        6          -           -     -           -
      -              N/A      00:57:35    5          -           Yes   24966        -
      No        10.12.96.4    -

```

(via rsvp tunnel-id 5001,

nhlfe_ix 4, label 24320)

```
Pl#show mpls ilm-table
```

```
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
```

```
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
```

```
S - SNMP, L - LDP, R - RSVP, C - CRLDP
```

```
B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
```

```
O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
```

```
P - SR Policy, U - unknown, UPStr - upstream
```

```
ILM-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- | |
|-----------------------------------|---------------|----------|----------|-----------|---------|-------|-------------------|
| Intf/VRF | Nexthop | | pri | Algo-Num | UpTime | UPStr | peers |
| L> 30.1.1.0/24 | Yes N/A | 4 | 24964 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | 00:58:02 | 1 | | | | |
| L> 10.12.96.4/32 | Yes N/A | 2 | 24962 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | 00:58:02 | 1 | | | | |
| L> 10.12.96.3/32 | Yes N/A | 1 | 24961 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | 00:58:02 | 1 | | | | |
| L> 10.12.96.5/32 | Yes - | 3 | 24963 | 24966 | N/A | N/A | 10.12.96.4 |
| | | - | | | | | |
| | | | | | | | (via rsvp tunnel- |
| id 5001, nhlfe_ix 4, label 24320) | | | | | | | |
| L> 10.12.96.4/32 | Yes N/A | 8 | 24968 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | 00:58:02 | 1 | | | | |
| L> 10.12.96.1/32 | Yes N/A | 6 | 24966 | 3 | N/A | xe20 | 10.1.1.1 |
| | | 00:58:02 | 1 | | | | |
| L> 40.1.1.0/24 | Yes N/A | 5 | 24965 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | 00:58:02 | 1 | | | | |
| L> 10.12.96.3/32 | Yes N/A | 7 | 24967 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | 00:58:02 | 1 | | | | |
| L> 30.1.1.0/24 | Yes N/A | 10 | 24970 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | 00:58:02 | 1 | | | | |
| L> 10.12.96.5/32 | Yes N/A | 9 | 24969 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | 00:58:02 | 1 | | | | |
| L> 40.1.1.0/24 | Yes N/A | 11 | 24971 | Nolabel | N/A | N/A | 127.0.0.1 |
| | | 00:58:02 | 1 | | | | |

Commands for LDP Tunneling over RSVP-TE

The LDP Tunneling over RSVP-TE introduces the following configuration commands.

ldp-tunneling

Use this command to enable LDP tunneling over RSVP trunk. When a specific RSVP trunk is enabled for tunneling, user traffic is tunneled using LDP LSP over RSVP LSP. If more than one trunk is enabled for tunneling LDP LSP, following trunk selection method is followed:

- If there are more than one trunk with same tunnel end-node, trunk with best metric (lower cost) is selected.
- If a destination FEC is reachable via more than one tunnel-endpoint, a tunnel-endpoint which is closer to destination is selected for tunneling.



Note: TLDP sessions should be manually established with RSVP tunnel end-nodes. Additionally, the 'advertise-labels' CLI must be explicitly configured to permit label advertisement over TLDP sessions.

Use `no` parameter of this command to disable tunneling from a trunk.

Command Syntax

```
ldp-tunneling
no ldp-tunneling
```

Parameters

None

Default

Disabled

Command Mode

RSVP-trunk mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to enable LDP tunneling over RSVP trunk:

```
OcNOS#configure terminal
OcNOS(config)#rsvp-trunk t2
OcNOS(config-trunk)#to 4.4.4.4
OcNOS(config-trunk)#ldp-tunneling
OcNOS(config-trunk)#commit
OcNOS(config-trunk)#end
```

prefer-tunnel-in-tunnel rsvp

Use this command for prioritizing RSVP trunk over LDP-LSP for forwarding LDP traffic. By default incoming LDP traffic is forwarded using LDP LSP. However when this CLI is configured and if RSVP trunk has been enabled for tunneling LDP LSP, user data (incoming LDP LSP) is tunneled over RSVP tunnels. If this CLI is not enabled and RSVP trunk has been enabled for tunneling LDP LSP, user data still can be forwarded over RSVP trunk if no LDP LSP exist.

Use `no` parameter of this command to prioritizing LDP-LSP over RSVP trunk while forwarding LDP traffic.

Command Syntax

```
prefer-tunnel-in-tunnel rsvp
no prefer-tunnel-in-tunnel rsvp
```

Parameters

None

Default

LDP-LSP is selected over RSVP trunks for forwarding.

Command Mode

Router LDP mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to prioritize RSVP trunk over LDP-LSP for forwarding LDP traffic:

```
OcNOS#configure terminal
OcNOS(config)#router ldp
OcNOS(config-router)#prefer-tunnel-in-tunnel rsvp
OcNOS(config-router)#commit
OcNOS(config-router)#end
```


show ldp tunneling

This command displays the LDP tunneling.

Command Syntax

```
show ldp tunneling
```

Parameters

None

Default

None

Command Mode

EXEC mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to view the LDP tunneling on router R2:

```
#show ldp tunneling
Tunnel Name       : t1
Tunnel Endpoint   : 10.12.96.4/32
Tunnel Cost       : 20
Tunnel Owner      : RSVP
Tunnel Status     : Up

FEC                Upstream-Peer  In-Label  Out-Label
=====
10.12.96.5/32      10.12.96.1  26242    26244

Total FEC tunneld by t1 : 1
```

show ldp tunneling fec

This command displays the LDP tunneling FEC mappings.

Command Syntax

```
show ldp tunneling-fec
```

Parameters

None

Default

None

Command Mode

EXEC mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following configuration illustrates how to view the FEC mappings on router R2:

```
R2#show ldp tunneling-fec
FEC          Tunnel-name      Tunnel-endpoint  Upstream-Peer  In-label  Out-
label
52.1.1.0/24   t2                          4.4.4.4/32      1.1.1.1        26253    26250
53.1.1.0/24   t2                          4.4.4.4/32      1.1.1.1        26255    26241

Total LDP Tunneled FEC : 2
```

show ldp tunneling-tunnels

This command displays the LDP tunneling tunnels.

Command Syntax

```
show ldp tunneling-tunnels
```

Parameters

None

Default

None

Command Mode

EXEC mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to view the LDP tunneling on router R2:

```
R2#show ldp tunneling-tunnels
Tunnel-name      Tunnel-endpoint  Status  Cost
t2               4.4.4.4/32      Up      20
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms or Acronym | Description |
|--------------------------------------|--|
| Forward Error Correction (FEC) | A system of error control that allows the receiver to correct some errors without having to request a re-transmission of data. |
| Interior Gateway Protocol (IGP) | An intradomain protocol used to exchange network reachability and routing information among devices within an autonomous system (AS), such as Intermediate System to Intermediate System (IS-IS), Open Shortest Path First (OSPF), or Routing Information Protocol (RIP). Contrast with Exterior Gateway Protocol (EGP). |
| Label Distribution Protocol (LDP) | A protocol for distributing labels in non-traffic-engineered applications. LDP allows routers to create label-switched path (LSP) instances through a network by mapping network layer routing information directly to data-link layer switched paths. |
| Resource Reservation Protocol (RSVP) | A signaling protocol for reserving resources across a network. RSVP is rarely used by itself, but Resource Reservation Protocol—Traffic Engineering (RSVP-TE) is widely used. |

| Key Terms or Acronym | Description |
|---|---|
| Targeted Label Distribution Protocol (TLDP) | A specialized form of LDP (Label Distribution Protocol) sessions. |

Revert Timer for Revertive Switchover for LDP Pseudowires

Overview

The Virtual Circuit (VC) Reversion functionality, along with an associated Revert Timer, enhances traffic convergence and service stability in OcNOS LDP-based Layer 2 VPN deployments. It ensures that in dual-homing or redundant configurations, the Primary VC is automatically restored and preferred once it becomes available after a fault condition is resolved.

Characteristics

- VC reversion operates only in Primary/Secondary VC configurations. In case of a fault on the Attachment Circuit (AC) or Packet Switched Network (PSN), the system may bring down either or both VCs. Once the fault is resolved, the Primary VC will automatically revert to the active state.
- For VPWS deployments, a revert timer has been introduced to allow controlled, delayed reactivation of the Primary VC after recovery.
- For VPLS and H-VPLS topologies, VC reversion functionality has been newly added, including support for the revert timer mechanism.

Benefits

- Ensures deterministic traffic recovery by prioritizing the Primary VC when restored.
- Minimizes traffic blackholing during transient network recovery events.

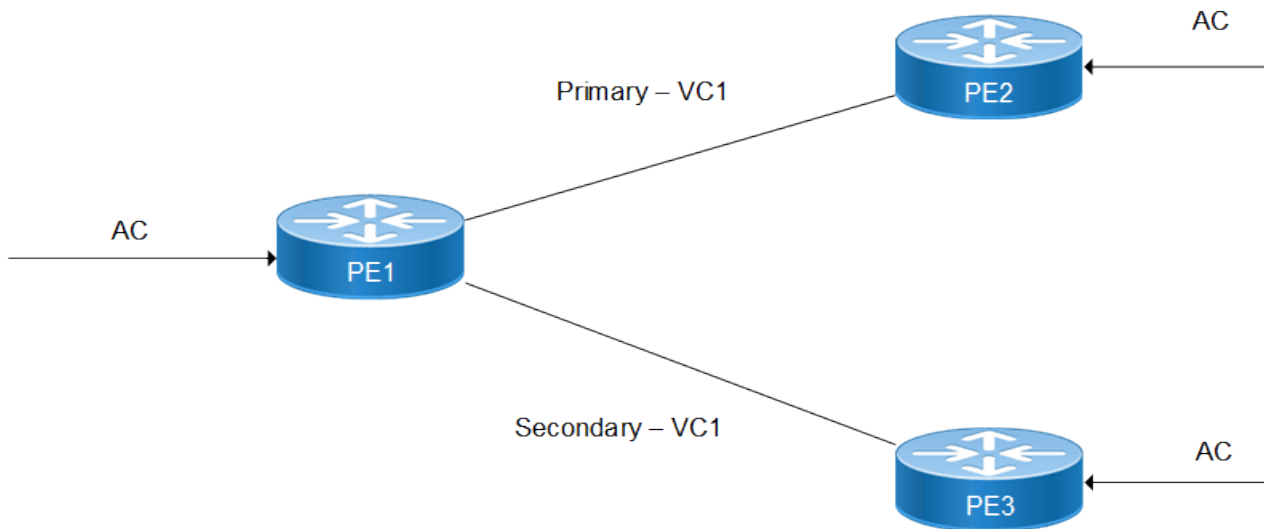
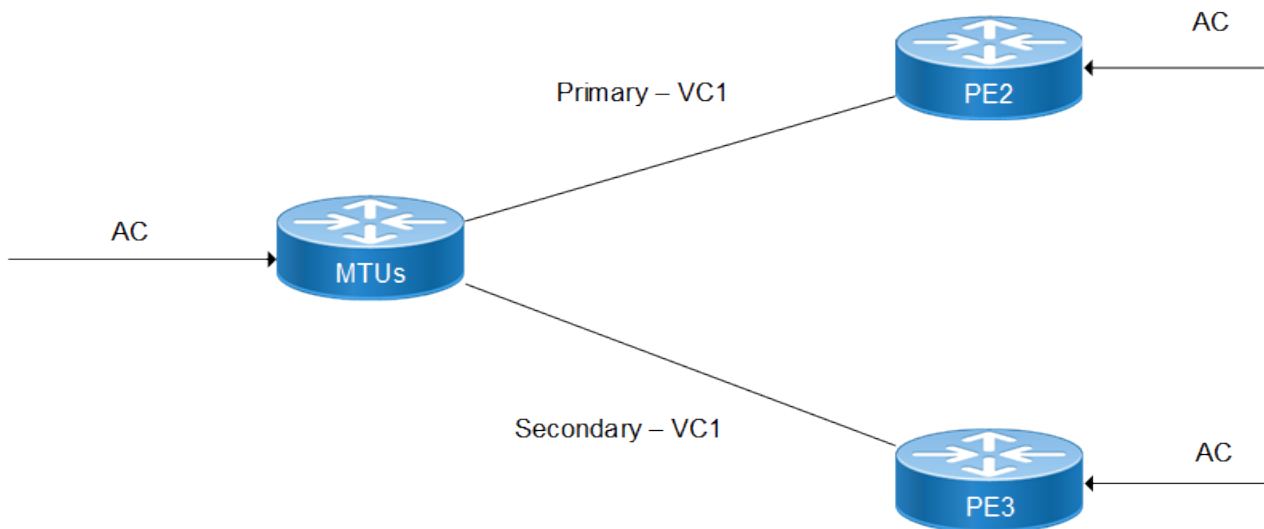
Prerequisites

- Applicable only to LDP-based VPWS and VPLS/H-VPLS services.
- Requires Primary/Secondary VC configuration for reversion logic to apply.

Configuration

Topology

In the following topology diagrams, faults may occur on the AC or the PSN at PE1, PE2, or PE3. Each fault triggers signaling that indicates whether the VC is in a forwarding or non-forwarding state. When the fault is cleared and the VC is restored, the Primary VC becomes eligible for reversion. If VC reversion is enabled, the switch back to the Primary VC occurs after the delay specified by the revert timer.

Figure 19. VPLS - VPWS VC Reversion Scenario**Figure 20. VPLS - H-VPLS VC Reversion scenario**

MPLS LDP Configuration

PE1

1. Enable LDP and define router ID.

```
(config)#router ldp
(config-router)#router-id 28.28.28.28
(config-router)#pw-status-tlv
(config-router)#commit
(config-router)#exit
```

2. Configure the interface.

```
(config)#interface eth0
(config-if)#ip vrf forwarding management
(config-if)#ip address dhcp
(config-if)#commit
(config-if)#exit
```

3. Configure VLAN sub-interface and IP address on ge0.

```
(config)#interface ge0
(config-if)#commit
(config-if)#exit

(config)#interface ge0.99
(config-if)#encapsulation dot1q 280
(config-if)#ip address 192.168.1.28/24
(config-if)#commit
(config-if)#exit
```

4. Configure VLAN sub-interface and IP address on ge1.

```
(config)#interface ge1
(config-if)#commit
(config-if)#exit

(config)#interface ge1.99
(config-if)#encapsulation dot1q 99
(config-if)#ip address 99.2.0.28/24
(config-if)#label-switching
(config-if)#enable-ldp ipv4
(config-if)#commit
(config-if)#exit
```

5. Configure the loopback interfaces.

```
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 28.28.28.28/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit

(config)#interface lo.management
(config-if)#ip vrf forwarding management
(config-if)#ip address 127.0.0.1/8
(config-if)#ipv6 address ::1/128
```

6. Configure the interfaces.

```
(config)#interface xe20
(config-if)#commit
(config-if)#exit

(config)#interface xe21
(config-if)#commit
(config-if)#exit
```

7. Configure VLAN sub-interface and IP address on xe21.

```
(config)#interface xe21.99
(config-if)#encapsulation dot1q 99
(config-if)#ip address 99.1.0.28/24
(config-if)#label-switching
(config-if)#enable-ldp ipv4
(config-if)#commit
(config-if)#exit
```

8. Add static routes.

```
(config)#ip route 29.29.29.29/32 99.2.0.29
(config)#ip route 32.32.32.32/32 99.5.0.32
(config)#commit
(config)#exit
```

PE2**1. Enable LDP and define router ID.**

```
(config)#router ldp
(config-router)#router-id 29.29.29.29
(config-router)#pw-status-tlv
(config-router)#commit
(config-router)#exit
```

2. Configure the interface.

```
(config)#interface eth0
(config-if)#ip vrf forwarding management
(config-if)#ip address dhcp
(config-if)#commit
(config-if)#exit
```

3. Configure VLAN sub-interface and IP address on ge1.

```
(config)#interface ge1.99
(config-if)#encapsulation dot1q 99
(config-if)#ip address 99.2.0.29/24
(config-if)#label-switching
(config-if)#enable-ldp ipv4
(config-if)#commit
(config-if)#exit
```

4. Configure VLAN sub-interface and IP address on ge2.

```
(config)#interface ge2.99
(config-if)#encapsulation dot1q 99
(config-if)#ip address 99.3.0.29/24
(config-if)#label-switching
(config-if)#enable-ldp ipv4
(config-if)#commit
(config-if)#exit
```

5. Configure the loopback interfaces.

```
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 29.29.29.29/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit
```

```
(config)#interface lo.management
(config-if)#ip vrf forwarding management
(config-if)#ip address 127.0.0.1/8
(config-if)#ipv6 address ::1/128
```

6. Add static routes.

```
(config)#ip route 28.28.28.28/32 99.2.0.28
(config)#ip route 32.32.32.32/32 99.3.0.32
(config)#commit
(config)#exit
```

PE3**1. Enable LDP and define router ID.**

```
(config)#router ldp
(config-router)#router-id 32.32.32.32
(config-router)#pw-status-tlv
(config-router)#commit
(config-router)#exit
```


2. Configure the interface.

```
(config)#interface eth0
(config-if)#ip vrf forwarding management
(config-if)#ip address dhcp
(config-if)#commit
(config-if)#exit
```

3. Configure VLAN sub-interface and IP address on ge0.

```
(config)#interface ge0.99
(config-if)#encapsulation dot1q 99
(config-if)#ip address 99.3.0.32/24
(config-if)#commit
(config-if)#exit
```

4. Configure VLAN sub-interface and IP address on ge3.

```
(config)#interface ge3.99
(config-if)#encapsulation dot1q 99
(config-if)#ip address 99.5.0.32/24
(config-if)#label-switching
(config-if)#enable-ldp ipv4
(config-if)#commit
(config-if)#exit
```

5. Configure the loopback interfaces.

```
(config)#interface lo
(config-if)#ip address 127.0.0.1/8
(config-if)#ip address 2.2.2.2/32 secondary
(config-if)#ip address 32.32.32.32/32 secondary
(config-if)#ipv6 address ::1/128
(config-if)#commit
(config-if)#exit
```

```
(config)#interface lo.management
(config-if)#ip vrf forwarding management
(config-if)#ip address 127.0.0.1/8
(config-if)#ipv6 address ::1/128
```

6. Add static routes.

```
(config)#ip route 28.28.28.28/32 99.5.0.28
(config)#ip route 29.29.29.29/32 99.3.0.29
(config)#commit
(config)#exit
```

VPWS Configuration

PE1

1. Configure the access interface.

```
(config)#interface ge2.21 switchport
(config-if)#encapsulation dot1q 282
(config-if)#access-if-vpws
(config-acc-if-vpws)#mpls-l2-circuit vc-2829 primary
(config-acc-if-vpws)#mpls-l2-circuit vc-2832 secondary
(config-acc-if-vpws)#vc-mode revertive timer 100
(config-acc-if-vpws)#exit
(config-if)#commit
(config-if)#exit
```

2. Define MPLS L2 circuits.

```
(config)#mpls l2-circuit vc-2829 2829 29.29.29.29
```

```
(config)#mpls l2-circuit vc-2832 2832 32.32.32.32
(config)#commit
(config)#exit
```

PE2

1. Configure the access interface.

```
(config)#interface ge0.21 switchport
(config-if)#encapsulation dot1q 282
(config-if)#access-if-vpws
(config-acc-if-vpws)#mpls-l2-circuit vc-2829 primary
(config-acc-if-vpws)#exit
(config-if)#commit
(config-if)#exit
```

2. Define MPLS L2 circuits.

```
(config)#mpls l2-circuit vc-2829 2829 28.28.28.28
(config)#commit
(config)#exit
```

PE3

1. Configure the access interface.

```
(config)#interface xe15.21 switchport
(config-if)#encapsulation dot1q 282
(config-if)#access-if-vpws
(config-acc-if-vpws)#mpls-l2-circuit vc-2832 primary
(config-acc-if-vpws)#exit
(config-if)#commit
(config-if)#exit
```

2. Define MPLS L2 circuits.

```
(config)#mpls l2-circuit vc-2832 2832 28.28.28.28
(config)#commit
(config)#exit
```

VPLS Configuration

PE1

1. Configure a VPLS instance with LDP signaling, defining primary and secondary virtual circuits (vc) in revertive mode to provide redundant Layer-2 VPN connectivity for a VLAN-based service.

```
(config)#mpls vpls vpls731 731
(config-vpls)#vpls-vc vc-28291
(config-vpls-spoke)#secondary vc-28321
(config-vpls-spoke)#vc-mode revertive timer 60
(config-vpls-spoke)#exit-spoke
(config-vpls)#signaling ldp
(config-vpls-sig)#vpls-type vlan
(config-vpls-sig)#ignore-ac-spoke-state
(config-vpls-sig)#exit-signaling
(config-vpls)#commit
(config-vpls)#exit-vpls
```

2. Configure the access interface.

```
(config)#interface ge2.99 switchport
(config-if)#encapsulation dot1q 99
(config-if)#access-if-vpls
(config-acc-if-vpls)#mpls-vpls vpls731
```

```
(config-acc-if-vpls)#exit
(config-if)#commit
(config-if)#exit
```

3. Define MPLS L2 circuits.

```
(config)#mpls l2-circuit vc-28291 28291 29.29.29.29
(config)#mpls l2-circuit vc-28321 28321 32.32.32.32
```

PE2

1. Configure a VPLS instance with LDP signaling, defining primary and secondary virtual circuits (vc) in revertive mode to provide redundant Layer-2 VPN connectivity for a VLAN-based service.

```
(config)#mpls vpls vpls731 731
(config-vpls)#vpls-vc vc-28291
(config-vpls-spoke)#exit-spoke
(config-vpls)#signaling ldp
(config-vpls-sig)#vpls-type vlan
(config-vpls-sig)#ignore-ac-spoke-state
(config-vpls-sig)#exit-signaling
(config-vpls)#commit
(config-vpls)#exit-vpls
```

2. Configure the access interface.

```
(config)#interface ge2.99 switchport
(config-if)#encapsulation dot1q 99
(config-if)#access-if-vpls
(config-acc-if-vpls)#mpls-vpls vpls731
(config-acc-if-vpls)#exit
(config-if)#commit
(config-if)#exit
```

3. Define MPLS L2 circuits.

```
(config)#mpls l2-circuit vc-28291 28291 29.29.29.29
(config)#mpls l2-circuit vc-28321 28321 32.32.32.32
```

PE3

1. Configure a VPLS instance with LDP signaling, defining primary and secondary virtual circuits (vc) in revertive mode to provide redundant Layer-2 VPN connectivity for a VLAN-based service.

```
(config)#mpls vpls vpls731 731
(config-vpls)#vpls-vc vc-28321
(config-vpls-spoke)#exit-spoke
(config-vpls)#signaling ldp
(config-vpls-sig)#vpls-type vlan
(config-vpls-sig)#ignore-ac-spoke-state
(config-vpls-sig)#exit-signaling
(config-vpls)#commit
(config-vpls)#exit-vpls
```

2. Configure the access interface.

```
(config)#interface xe13.99 switchport
(config-if)#encapsulation dot1q 99
(config-if)#access-if-vpls
(config-acc-if-vpls)#mpls-vpls vpls731
(config-acc-if-vpls)#exit
(config-if)#commit
(config-if)#exit
```

3. Define MPLS L2 circuit.

```
mpls l2-circuit vc-28321 28321 28.28.28.28
```

Reverter Timer Running Configuration

MPLS LDP Configuration

PE1

```
!  
router ldp  
  router-id 28.28.28.28  
  pw-status-tlv  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface ge0  
!  
interface ge0.99  
  encapsulation dot1q 280  
  ip address 192.168.1.28/24  
!  
interface ge1  
!  
interface ge1.99  
  encapsulation dot1q 99  
  ip address 99.2.0.28/24  
  label-switching  
  enable-ldp ipv4  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 28.28.28.28/32 secondary  
  ipv6 address ::1/128  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xe20  
!  
interface xe21  
!  
interface xe21.99  
  encapsulation dot1q 99  
  ip address 99.1.0.28/24  
  label-switching  
  enable-ldp ipv4  
!  
ip route 29.29.29.29/32 99.2.0.29  
ip route 32.32.32.32/32 99.5.0.32  
!  
exit  
!  
end
```

PE2

```
!  
router ldp  
  router-id 29.29.29.29  
  pw-status-tlv  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp
```

```
!  
interface ge1.99  
  encapsulation dot1q 99  
  ip address 99.2.0.29/24  
  label-switching  
  enable-ldp ipv4  
!  
interface ge2.99  
  encapsulation dot1q 99  
  ip address 99.3.0.29/24  
  label-switching  
  enable-ldp ipv4  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 29.29.29.29/32 secondary  
  ipv6 address ::1/128  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128!  
!  
ip route 28.28.28.28/32 99.2.0.28  
ip route 32.32.32.32/32 99.3.0.32  
!  
exit  
!  
end
```

PE3

```
!  
router ldp  
  router-id 32.32.32.32  
  pw-status-tlv  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface ge0.99  
  encapsulation dot1q 99  
  ip address 99.3.0.32/24  
  label-switching  
  enable-ldp ipv4  
!  
interface ge3.99  
  encapsulation dot1q 99  
  ip address 99.5.0.32/24  
  label-switching  
  enable-ldp ipv4  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 2.2.2.2/32 secondary  
  ip address 32.32.32.32/32 secondary  
  ipv6 address ::1/128  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
ip route 28.28.28.28/32 99.5.0.28  
ip route 29.29.29.29/32 99.3.0.29  
!  
exit
```

```
!  
end
```

VPWS Configuration

PE1

```
!  
interface ge2.21 switchport  
  encapsulation dot1q 282  
  access-if-vpws  
    mpls-l2-circuit vc-2829 primary  
    mpls-l2-circuit vc-2832 secondary  
    vc-mode revertive timer 100  
!  
mpls l2-circuit vc-2829 2829 29.29.29.29  
!  
mpls l2-circuit vc-2832 2832 32.32.32.32  
!  
exit  
!  
end
```

PE2

```
!  
interface ge0.21 switchport  
  encapsulation dot1q 282  
  access-if-vpws  
    mpls-l2-circuit vc-2829 primary  
!  
mpls l2-circuit vc-2829 2829 28.28.28.28  
!
```

PE3

```
!  
interface xe15.21 switchport  
  encapsulation dot1q 282  
  access-if-vpws  
    mpls-l2-circuit vc-2832 primary  
!  
mpls l2-circuit vc-2832 2832 28.28.28.28  
!  
exit  
!  
end
```

VPLS Configuration

PE1

```
!  
mpls vpls vpls731 731  
  vpls-vc vc-28291  
    secondary vc-28321  
    vc-mode revertive timer 60  
  exit-spoke  
  signaling ldp  
  vpls-type vlan  
  ignore-ac-spoke-state
```

```
exit-signaling
exit-vpls
!
interface ge2.99 switchport
encapsulation dot1q 99
access-if-vpls
mpls-vpls vpls731
!
mpls 12-circuit vc-28291 28291 29.29.29.29
!
mpls 12-circuit vc-28321 28321 32.32.32.32
!
exit
!
end
```

PE2

```
!
mpls vpls vpls731 731
vpls-vc vc-28291
exit-spoke
signaling ldp
vpls-type vlan
ignore-ac-spoke-state
exit-signaling
exit-vpls
!
interface ge0.20 switchport
encapsulation dot1q 280
access-if-vpls
mpls-vpls vpls73
!
mpls 12-circuit vc-28291 28291 28.28.28.28
!
exit
!
end
```

PE3

```
!
mpls vpls vpls731 731
vpls-vc vc-28321
exit-spoke
signaling ldp
vpls-type vlan
exit-signaling
exit-vpls
!
interface xe13.99 switchport
encapsulation dot1q 99
access-if-vpls
mpls-vpls vpls731
!
mpls 12-circuit vc-28321 28321 28.28.28.28
!
exit
!
end
```

Validation

VPWS Verification

VC Mode Revertive

```
PE1#show mpls l2-circuit detail
MPLS Layer-2 Virtual Circuit: vc-2829, id: 2829 PW-INDEX: 4 service-tpid: dot1.
q
Endpoint: 29.29.29.29
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
In-label: 24964, Out-label: 24962
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: ge2.21
  Subinterface Match Criteria(s) :
    dot1q 282

Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Revertive
Revertive immediate
Local PW Status :
  Forwarding
Remote PW Status :
  Forwarding
Local VCCV Capability:
  CC-Types: Type 2(in use)
  CV-Types: LSP ping(in use)

Virtual Circuit is active (State: Installed)
  CreateTime: 01:07:20, UpTime: 00:02:41

MPLS Layer-2 Virtual Circuit: vc-2832, id: 2832 PW-INDEX: 3 service-tpid: dot1.
q
Endpoint: 32.32.32.32
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
In-label: 24963, Out-label: 24963
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: ge2.21
  Subinterface Match Criteria(s) :
    dot1q 282

Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Secondary
Virtual Circuit is configured as Revertive
Revertive immediate
Local PW Status :
  PW Standby
Remote PW Status :
  Forwarding
Local VCCV Capability:
  CC-Types: Type 2(in use)
  CV-Types: LSP ping(in use)

Virtual Circuit is inactive (State: Signaled)
  CreateTime: 01:07:41, DownReason: Primary entity UP
```

VC Mode Revertive with Timer Running

```
PE1
#show mpls l2-circuit detail
MPLS Layer-2 Virtual Circuit: vc-2829, id: 2829 PW-INDEX: 4 service-tpid: dot1.
q
```



```

Endpoint: 29.29.29.29
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
In-label: 24964, Out-label: 24962
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: ge2.21
  Subinterface Match Criteria(s) :
    dot1q 282

Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Revertive
Revertive time 100 seconds, timer started, remaining 93 seconds
Local PW Status :
  PW Standby
Remote PW Status :
  Forwarding
Local VCCV Capability:
  CC-Types: Type 2(in use)
  CV-Types: LSP ping(in use)

Virtual Circuit is inactive (State: Signaled)
  CreateTime: 01:03:05, DownReason: VC on standby (non-revertive)

MPLS Layer-2 Virtual Circuit: vc-2832, id: 2832  PW-INDEX: 3 service-tpid: dot1.
q
Endpoint: 32.32.32.32
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
In-label: 24963, Out-label: 24963
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: ge2.21
  Subinterface Match Criteria(s) :
    dot1q 282

Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Secondary
Virtual Circuit is configured as Revertive
Revertive time 100 seconds
Local PW Status :
  Forwarding
Remote PW Status :
  Forwarding
Local VCCV Capability:
  CC-Types: Type 2(in use)
  CV-Types: LSP ping(in use)

Virtual Circuit is active (State: Installed)
  CreateTime: 01:03:26, UpTime: 01:01:19

```

VC Mode Revertive with Timer not Running

```

#show mpls l2-circuit detail
MPLS Layer-2 Virtual Circuit: vc-2829, id: 2829  PW-INDEX: 4 service-tpid: dot1.
q
Endpoint: 29.29.29.29
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
In-label: 24964, Out-label: 24962
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: ge2.21
  Subinterface Match Criteria(s) :
    dot1q 282

Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Revertive

```

```

Revertive time 100 seconds
Local PW Status :
  Forwarding
Remote PW Status :
  Forwarding
Local VCCV Capability:
  CC-Types:  Type 2(in use)
  CV-Types:  LSP ping(in use)

Virtual Circuit is active  (State: Installed)
  CreateTime: 01:04:42, UpTime: 00:00:03

MPLS Layer-2 Virtual Circuit: vc-2832, id: 2832  PW-INDEX: 3 service-tpid: dot1.
q
Endpoint: 32.32.32.32
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
In-label: 24963, Out-label: 24963
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: ge2.21
  Subinterface Match Criteria(s) :
    dot1q 282

Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Secondary
Virtual Circuit is configured as Revertive
Revertive time 100 seconds
Local PW Status :
  PW Standby
Remote PW Status :
  Forwarding
Local VCCV Capability:
  CC-Types:  Type 2(in use)
  CV-Types:  LSP ping(in use)

Virtual Circuit is inactive (State: Signaled)
  CreateTime: 01:05:03, DownReason: Primary entity UP

```

VPLS Verification

VC Mode Revertive

```

PE1#show mpls vpls detail
Virtual Private LAN Service Instance: vpls731, ID: 731
  SIG-Protocol: LDP
  Attachment-Circuit: UP
  Learning: Enabled
  Control-Word: Disabled
  Flow Label Status: Disabled, Direction: None, Static: No
  Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
  Description: none
  service-tpid: dot1.q
  Operating mode: Tagged
  Svlan Id: 0
  Svlan Tpid: 8100
  Ignoring AC interface and spoke-VC state
  MAC Withdrawal:

Configured interfaces:
  Interface: ge2.99
    Status: Up
    Subinterface Match Criteria(s) :
      dot1q 99

Spoke Peers:

```

```
vc-28291 (Up) (UpTime 01:01:31)
Secondary: vc-28321 (Dn) (Reason: VC on standby)
vc-mode : Revertive, Revertive immediate
```

VC Mode Revertive with Timer Running

```
PE1#show mpls vpls detail
Virtual Private LAN Service Instance: vpls731, ID: 731
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
Ignoring AC interface and spoke-VC state
MAC Withdrawal:

Configured interfaces:
Interface: ge2.99
Status: Up
Subinterface Match Criteria(s) :
dot1q 99

Spoke Peers:
vc-28291 (Dn) (Reason: VC on standby)
Secondary: vc-28321 (Up) (UpTime 00:00:54)
vc-mode : Revertive, Revertive time 60 seconds
Timer started, remaining 12 seconds
```

VC Mode Revertive with Timer not Running

```
PE1#show mpls vpls detail
Virtual Private LAN Service Instance: vpls731, ID: 731
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
Ignoring AC interface and spoke-VC state
MAC Withdrawal:

Configured interfaces:
Interface: ge2.99
Status: Up
Subinterface Match Criteria(s) :
dot1q 99

Spoke Peers:
vc-28291 (Up) (UpTime 01:02:37)
Secondary: vc-28321 (Dn) (Reason: VC on standby)
vc-mode : Revertive, Revertive time 60 seconds
```

Revert Timer Commands

The following commands are revised as part of the VC reversion functionality.

- [vc-mode \(page 1624\)](#)

Implementation Examples

Scenario 1: In an LDP-based VPWS or VPLS/H-VPLS deployment, Primary and Secondary pseudowires are configured for redundancy. When the Primary VC fails due to an AC or PSN issue, traffic switches to the Secondary VC. After the fault is cleared, immediate reversion to the Primary VC can cause traffic flapping during unstable recovery.

Use Case 1: Implementing revertive VC switchover with a configurable revert timer is essential to fulfill this requirement. When the Primary VC recovers, traffic continues on the Secondary VC until the revert timer expires. If the Primary VC remains stable for the configured duration, traffic switches back to it automatically.

Troubleshooting Revert Timer for Revertive Switchover Issue

1. Primary VC does not revert after recovery
 - Revertive mode is not enabled
 - Revert timer is not configured
 - Primary VC is still in non-forwarding state
2. Revert timer not running
 - Primary VC never transitioned to DOWN state
 - VC is configured as non-revertive
 - Immediate revertive mode is in effect

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms/Acronym | Description |
|-------------------|------------------------------|
| H-VPLS | Hierarchical VPLS |
| LDP | Label Distribution Protocol |
| VPWS | Virtual Private Wire Service |
| VPLS | Virtual Private LAN Service |

LABEL DISTRIBUTION PROTOCOL COMMAND REFERENCE

| | |
|---|-----|
| LDP Commands | 516 |
| auth md5 password | 518 |
| advertise-labels | 519 |
| advertise-label-for-default-route | 521 |
| advertisement-mode | 522 |
| auto-targeted-session | 523 |
| clear ldp adjacency | 524 |
| clear ldp session | 525 |
| clear ldp statistics | 526 |
| clear ldp statistics advertise-labels | 527 |
| control-mode | 528 |
| debug ldp advertise-labels | 530 |
| debug ldp all | 531 |
| debug ldp dsm | 532 |
| debug ldp events | 533 |
| debug ldp fsm | 534 |
| debug ldp hexdump | 535 |
| debug ldp inter-area | 536 |
| debug ldp nsm | 537 |
| debug ldp packet | 538 |
| debug ldp usm | 539 |
| debug ldp vc usm | 540 |
| disable-ldp | 541 |
| enable-ldp | 542 |
| explicit-null | 543 |
| fast-reroute | 544 |
| global-merge-capability | 545 |
| graceful-restart | 546 |
| hello-interval | 548 |
| hold-time | 549 |
| import-bgp-routes | 550 |
| inter-area-lsp | 551 |
| keepalive-interval | 552 |
| keepalive-timeout | 553 |
| label-retention-mode | 554 |

| | |
|-------------------------------------|-----|
| ldp advertisement-mode | 556 |
| ldp hello-interval | 557 |
| ldp hold-time | 558 |
| ldp keepalive-interval | 559 |
| ldp keepalive-timeout | 560 |
| ldp label-retention-mode | 561 |
| ldp multicast-hellos | 562 |
| ldp-optimization | 563 |
| loop-detection | 564 |
| loop-detection-hop-count | 565 |
| loop-detection-path-vec-count | 566 |
| mpls ldp-igp sync-delay | 567 |
| mpls ldp-igp sync isis | 568 |
| mpls ldp-igp sync ospf | 569 |
| multicast-hellos | 570 |
| neighbor | 571 |
| neighbor auth md5 password | 572 |
| neighbor prefix-list | 573 |
| propagate-release | 574 |
| pw-status-tlv | 575 |
| request-labels-for | 576 |
| request-retry | 577 |
| request-retry-timeout | 578 |
| restart ldp graceful | 579 |
| rlfa-ilm-optimization | 580 |
| router-id | 581 |
| router ldp | 582 |
| session-group name | 583 |
| snmp restart ldp | 584 |
| targeted-peer hello-interval | 585 |
| targeted-peer ipv4 | 586 |
| transport-address ipv4 | 587 |
| targeted-peer tunneling | 588 |
| targeted-peer hold-time | 589 |
| targeted-peer-hold-time | 590 |
| LDP Show Commands | 591 |
| show ldp label-pool | 592 |
| show debugging ldp | 593 |
| show ldp | 594 |
| show ldp adjacency | 596 |

| | |
|--|-----|
| show ldp adjacency count | 597 |
| show ldp advertise-labels | 598 |
| show ldp fec | 599 |
| show ldp igp sync | 601 |
| show ldp inter-area-fecs | 602 |
| show ldp inter-area-fecs prefix | 603 |
| show ldp interface | 604 |
| show ldp lsp | 606 |
| show ldp mpls-l2-circuit | 609 |
| show ldp rfa-routes | 612 |
| show ldp rfa-routes count | 613 |
| show ldp routes | 614 |
| show ldp routes count | 616 |
| show ldp session | 617 |
| show ldp session all count | 620 |
| show ldp session multicast count | 621 |
| show ldp session targeted count | 622 |
| show ldp statistics | 623 |
| show ldp statistics advertise-labels | 625 |
| show ldp targeted-peers | 626 |
| show ldp upstream | 627 |
| show ldp downstream | 629 |
| show ldp vpls | 631 |
| show mpls ldp discovery | 634 |
| show mpls ldp neighbor | 635 |
| show mpls ldp parameter | 636 |
| show ldp tunneling-fec | 638 |
| show ldp tunneling | 639 |
| show ldp tunneling-tunnels | 640 |

LDP Commands

This chapter is a reference for the LDP commands:

| | |
|---|-----|
| auth md5 password | 518 |
| advertise-labels | 519 |
| advertise-label-for-default-route | 521 |
| advertisement-mode | 522 |
| auto-targeted-session | 523 |
| clear ldp adjacency | 524 |
| clear ldp session | 525 |
| clear ldp statistics | 526 |
| clear ldp statistics advertise-labels | 527 |
| control-mode | 528 |
| debug ldp advertise-labels | 530 |
| debug ldp all | 531 |
| debug ldp dsm | 532 |
| debug ldp events | 533 |
| debug ldp fsm | 534 |
| debug ldp hexdump | 535 |
| debug ldp inter-area | 536 |
| debug ldp nsm | 537 |
| debug ldp packet | 538 |
| debug ldp usm | 539 |
| debug ldp vc usm | 540 |
| disable-ldp | 541 |
| enable-ldp | 542 |
| explicit-null | 543 |
| fast-reroute | 544 |
| global-merge-capability | 545 |
| graceful-restart | 546 |
| hello-interval | 548 |
| hold-time | 549 |
| import-bgp-routes | 550 |
| inter-area-lsp | 551 |
| keepalive-interval | 552 |
| keepalive-timeout | 553 |
| label-retention-mode | 554 |
| ldp advertisement-mode | 556 |
| ldp hello-interval | 557 |

| | |
|-------------------------------------|-----|
| ldp hold-time | 558 |
| ldp keepalive-interval | 559 |
| ldp keepalive-timeout | 560 |
| ldp label-retention-mode | 561 |
| ldp multicast-hellos | 562 |
| ldp-optimization | 563 |
| loop-detection | 564 |
| loop-detection-hop-count | 565 |
| loop-detection-path-vec-count | 566 |
| mpls ldp-igp sync-delay | 567 |
| mpls ldp-igp sync isis | 568 |
| mpls ldp-igp sync ospf | 569 |
| multicast-hellos | 570 |
| neighbor | 571 |
| neighbor auth md5 password | 572 |
| neighbor prefix-list | 573 |
| propagate-release | 574 |
| pw-status-tlv | 575 |
| request-labels-for | 576 |
| request-retry | 577 |
| request-retry-timeout | 578 |
| restart ldp graceful | 579 |
| rlfa-ilm-optimization | 580 |
| router-id | 581 |
| router ldp | 582 |
| session-group name | 583 |
| snmp restart ldp | 584 |
| targeted-peer hello-interval | 585 |
| targeted-peer ipv4 | 586 |
| transport-address ipv4 | 587 |
| targeted-peer tunneling | 588 |
| targeted-peer hold-time | 589 |
| targeted-peer-hold-time | 590 |

auth md5 password

Use this command to configure session-group authentication.

Command Syntax

```
(auth) md5 password (plain-text|encrypt) WORD
```

Parameters

password plain-text

Enter the password in plain text. This option stores the password as clear text, making it easily readable.

password encrypt

Use an encrypted password. This option encrypts the password while storing it, preventing unauthorized access.

WORD

Specify the actual password (either in plain text or encrypted).

Default

None

Command Mode

Session group mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#session-group name abcd
(config-router-sg)#auth md5 password plain-text test2
```

advertise-labels

Use this command to restrict or prevent the advertisement of locally assigned labels to LDP neighbors.

Use the `no` parameter to restore the default behavior, which advertises all locally assigned labels to all LDP neighbors.

Command Syntax

```
advertise-labels for any to none
advertise-labels for PREFIX to (PEER|any)
no advertise-labels for any to none
no advertise-labels for PREFIX to (PEER|any)
```

Parameters

for

Introduces the prefix list that defines which prefixes (destinations) will have labels advertised.

any

Matches all entries. When used with `for`, it selects all prefixes. When used with `to`, it selects all LDP peers.

PREFIX

Name of a configured IP prefix-list that specifies the set of prefixes (destinations) to advertise labels for.

to

Introduces the peer list that defines the neighbors receiving the advertisements.

PEER

Name of a configured IP prefix-list that specifies the LDP neighbors that receive the label advertisements.

none

Matches no entries. When used with `to`, it prevents label advertisement to all LDP peers.



Note: Both `PREFIX` and `PEER` must reference existing prefix-lists.

Default

OcNOS advertises labels for all prefixes to all LDP neighbors.

Command Mode

Router mode

Applicability

Introduced before OcNOS version 1.3.

Examples

Example 1: Advertise specific prefixes to specific peers

In this example:

- The `mylist` prefix-list defines the prefix(es) to advertise.
- The `peerlist` prefix-list defines the neighbors that should receive these advertisements.

```
#configure terminal
(config)#router ldp
(config-router)#advertise-labels for mylist to peerlist

!
ip prefix-list mylist
  seq 5 permit 14.1.0.70/32
  seq 10 deny 0.0.0.0/0
!
ip prefix-list peerlist
  seq 5 permit 14.1.0.0/24
  seq 10 deny 0.0.0.0/0
!
router ldp
  router-id 14.1.0.70
  targeted-peer ipv4 14.1.0.30
  exit-targeted-peer-mode
  transport-address ipv4 14.1.0.70
  advertise-labels for mylist to peerlist
```

Example 2: Prevent all label advertisements

```
#configure terminal
(config)#router ldp
(config-router)#advertise-labels for any to none
```

Example 3: Restore default (advertise all labels to all peers)

```
#configure terminal
(config)#router ldp
(config-router)#no advertise-labels for any to none
```

advertise-label-for-default-route

Use this command to enable label advertisement for default route.

Use `no` form to disable the label advertisement for default route.

Command Syntax

```
advertise-label-for-default-route
```

Parameters

None

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 4.2.

Examples

```
#configure terminal
(config)#router ldp
(config-router)#advertise-label-for-default-route
```

advertisement-mode

Use this command to set the label advertisement mode for all the interfaces for the current LSR. Specifying `downstream-on-demand` and `downstream-unsolicited` mode affects which LSR initiates mapping requests and mapping advertisements.

This command is a global command used to set the label advertisement mode for all interfaces for the current LSR. The advertisement mode set for a specific interface overrides the value set by this command (see `ldp advertisement-mode`). Use this command before starting the interface as it closes and restarts all sessions.

Use the `no` parameter to revert to the default advertisement mode value.

Command Syntax

```
advertisement-mode (downstream-on-demand|downstream-unsolicited)
no advertisement-mode (downstream-on-demand|downstream-unsolicited)
```

Parameters

downstream-on-demand

Sends label upon request. When a users uses this mode, a router distributes a label to a peer only if there is a pending label request from a peer. The reaction of the downstream router to this request depends on the label advertising mode supported on the next hop. This mode is typically used with the conservative label retention mode.

downstream-unsolicited

Sends label without waiting request. This mode distributes labels to peers without waiting for a label request, and is typically used with the liberal label retention mode.

Default

By default, advertisement mode is `downstream-unsolicited`

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

In the following example, the LSR will use the `downstream-unsolicited` advertisement mode for an LDP session on its interfaces.

```
#configure terminal
(config)#router ldp
(config-router)#advertisement-mode downstream-unsolicited
```

auto-targeted-session

Use this command on the remote LFA node to accept received targeted hello messages and enable dynamic creation of TLDP sessions for establishing rLFA tunnels. This command is only for remote LFA.

Use the `no` parameter to disable dynamic TLDP session creation.



Notes:

- Dynamically created TLDP sessions can only be deleted when "auto-targeted-session" or "fast-reroute" configurations are removed.
- The `auto-targeted-session` command advertises the labels to static TLDP sessions. To avoid advertising the labels to static TLDP sessions, configure the [advertise-labels \(page 519\)](#) command to prevent the label advertisement to static TLDP sessions.
- When the `auto-targeted-session` command establishes RLFA tunnels, a unique label is advertised to distinguish the primary LSP from the backup LSP. This ensures the forwarding plane can clearly differentiate between the active and backup paths.
- Remote LFA (RLFA) over an RLFA backup is not supported.

Command Syntax

```
auto-targeted-session
no auto-targeted-session
```

Parameters

None

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#configure terminal
(config)#router ldp
(config-router)#auto-targeted-session
(config-router)#
```

clear ldp adjacency

Use this command to clear an adjacency with a specified peer, or to clear all adjacencies for the current LSR.

Command Syntax

```
clear ldp adjacency (A.B.C.D|*)
```

Parameters

*

Specify to clear all adjacencies.

A.B.C.D

Specify to clear IPv4 address of the peer.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ldp adjacency 123.123.123.33
```


clear ldp session

Use this command to clear a session established with a specified peer, or to clear all sessions for the current LSR.

Command Syntax

```
clear ldp session (A.B.C.D|*)
```

Parameters

*

Specify to clear all sessions.

A.B.C.D

Specify to clear IPv4 address of the peer.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ldp session 123.123.123.33
```

clear ldp statistics

Use this command to clear LDP statistics. This command clears the count per each operation filtered by an advertisement list.

Command Syntax

```
clear ldp statistics
```

Parameters

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ldp statistics
```

clear ldp statistics advertise-labels

Use this command to clear LDP advertise-labels statistics. This command clears the count per each operation filtered by an advertisement list.

Command Syntax

```
clear ldp statistics advertise-labels
clear ldp statistics advertise-labels for PREFIX
clear ldp statistics advertise-labels for PREFIX to PEER
```

Parameters

advertise-labels

Specify the IP prefix list of advertise-labels.

for

Specify the permitted destinations.

PREFIX

Specify the destinations that have their labels advertised.

to

Specify the given neighbor.

PEER

Specify the LDP neighbors that receive these advertisements.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear ldp statistics advertise-labels
```

control-mode

Use this command to set the control mode for label processing. Ordered processing sets the mode to strict chain-of-command; an LSR replies to a request packet from an LSR higher in the chain only after it receives a label from an LSR lower in the chain. Independent processing sets the mode to instant replies.

In independent control mode, each LSR might advertise label mappings to its neighbors at any time. In independent downstream-on-demand mode, an LSR might answer requests for label mappings immediately, without waiting for a label mapping from the next hop. In independent downstream unsolicited mode, an LSR might advertise a label mapping for an Forwarding Equivalence Class (FEC) to its neighbors whenever it is prepared to label-switch that FEC. In independent mode, an upstream label can be advertised before a downstream label is received.

In ordered control mode, an LSR may initiate the transmission of label mapping only for an FEC for which it has a label mapping for the FEC next hop, or for which the LSR is the egress. For each FEC for which the LSR is not the egress and no mapping exists, the LSR must wait until a label from a downstream LSR is received. An LSR may be an egress for some FECs and a non-egress for others. Changes in control mode only affect labels that were sent or received after the change was made.

Use the `no` parameter to revert to default control mode.



Notes:

- Control mode "independent" is supported with advertisement mode "DU" only.
- When the advertisement mode is set as "DU", control mode automatically sets to "independent".
- Control mode "independent" is not supported with advertisement mode "DOD".
- Control mode "ordered" is supported with advertisement mode "DOD" only.
- Control mode "ordered" is not supported with advertisement mode "DU".
- When the advertisement mode is set as "DOD", control mode automatically sets to "ordered".

Command Syntax

```
control-mode (ordered|independent)
no control-mode
```

Parameters

independent

Sets control mode to independent processing.

ordered

Sets control mode to ordered processing.

Command Mode

Router mode

Default

Independent

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#control-mode ordered
```

debug ldp advertise-labels

Use this command to enable the debugging of LDP advertise-label events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp advertise-labels
no debug ldp advertise-labels
```

Parameters

None

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ldp advertise-labels
```

debug ldp all

Use this command to enable the debugging of all LDP events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp all
no debug ldp all
no debug all
```

Parameters

None

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ldp all
```

debug ldp dsm

Use this command to enable the debugging of LDP DSM events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp dsm
no debug ldp dsm
```

Parameters

None

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ldp dsm
```

debug ldp events

Use this command to enable the debugging of all LDP events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp events
no debug ldp events
```

Parameters

None

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ldp advertise-labels
(config)#debug ldp all
(config)#debug ldp dsm
(config)#debug ldp events
```

debug ldp fsm

Use this command to enable the debugging of LDP FSM events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp fsm
no debug ldp fsm
```

Parameters

None

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ldp fsm
```

debug ldp hexdump

Use this command to enable the debugging of LDP hexdump events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
targeted-peer-hold-time <3-65535>  
no targeted-peer-hold-time
```

Parameters

None

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#debug ldp hexdump
```

debug ldp inter-area

Use this command to enable the debugging of LDP inter-area events. On using the debug command, the router continues to generate an output until the no parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp inter-area
no debug ldp inter-area
```

Parameters

None

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 4.0.

Example

```
#configure terminal
(config)#debug ldp inter-area
```

debug ldp nsm

Use this command to enable the debugging of LDP NSM events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp nsm
no debug ldp nsm
```

Parameters

None

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ldp nsm
```

debug ldp packet

Use this command to enable the debugging of LDP packet events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp packet
debug ldp packet (notification|hello|initialization|keepalive|address|label)
no debug ldp packet
no debug ldp packet (notification|hello|initialization|keepalive|address|label)
```

Parameters

notification

Debug LDP notification packets.

hello

Debug LDP hello packets.

initialization

Debug LDP initialization packets.

keepalive

Debug LDP keepalive packets.

address

Debug LDP address (withdraw) packets.

label

Debug LDP address label packets.

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ldp packet hello
```

debug ldp usm

Use this command to enable the debugging of LDP USM events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp usm
no debug ldp usm
```

Parameters

None

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ldp usm
```

debug ldp vc usm

Use this command to enable the debugging of LDP VC events.

On using the debug command, the router continues to generate an output until the `no` parameter is used with this command. The debug output and system error messages are written on the virtual terminal.

Use the `no` parameter with this command to disable this function.

Command Syntax

```
debug ldp vc dsm
debug ldp vc usm
no debug ldp vc dsm
no debug ldp vc usm
```

Parameters

dsm

Debug LDP downstream SM.

usm

Debug LDP upstream SM.

Command Mode

Configure mode, Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#debug ldp vc dsm
(config)#debug ldp vc usm
```

disable-ldp

Use this command to disable LDP IPv4 on a specified interface.

This command disables the transmission of Hello packets through the current interface, and clears all created sessions and adjacencies for this interface. Use disable-ldp alone to disable only LDP IPv4 on the interface.

Command Syntax

```
disable-ldp (ipv4|)
```

Parameters

ipv4

Disables IPv4 on the interface.

Default

By default, disable ldp is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example disables LDP IPv4 on interface eth0.

```
#configure terminal
(config)#interface eth0
(config-if)#disable-ldp
```

The following example disables LDP IPv4 on interface eth0.

```
#configure terminal
(config)#interface eth0
(config-if)#disable-ldp ipv4
```

enable-ldp

Use this command to enable LDP IPv4 on a specified interface. This command enables the transmission of Hello packets through the current interface, so that LDP adjacencies and LDP sessions can be created.



Note: The corresponding interface must be enabled for label-switching using the [label-switching \(page 1396\)](#) command.

Command Syntax

```
enable-ldp ipv4
```

Parameters

None

Default

By default, enable ldp is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following example enables LDP IPv4 on interface eth0.

```
#configure terminal
(config)#interface eth0
(config-if)#enable-ldp ipv4
```

explicit-null

Use this command to configure the router to send explicit-null labels for directly connected FECs instead of implicit-null labels. Implicit-nulls are the default labels.

This command controls the label value advertised on the egress router of an LSP. By default, implicit null label (label 3) is advertised for directly connected FECs. LDP advertises an Implicit Null label that causes the previous hop router to perform penultimate hop popping. Use the `explicit null` command to avoid the penultimate router from penultimate hop popping, and to force it to replace the incoming label with the explicit null label.



Note: Do not use this command if the LDP is concurrently used for MPLS/BGP VPNs.

Use the `no` parameter to stop sending explicit-null labels for directly connected FECs and resume sending implicit-null labels for them.

Command Syntax

```
explicit-null
no explicit-null
```

Parameters

None

Default

By default, sends implicit-null labels.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ldp
(config-router)#explicit-null
```

fast-reroute

Use this command to enable the installation of backup paths advertised by IGP LFA (and/or rLFA) fast reroute.

Use no form CLI to disable the feature.

Command Syntax

```
fast-reroute
no fast-reroute
```

Parameters

None

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 5.1.

Examples

```
#configure terminal
(config)#router ldp
(config)#fast-reroute
```

global-merge-capability

Use this command to override the default merge capability setting of all the interfaces for the current LSR.

The merge capability aggregates multiple incoming flows with the same destination address into a single outgoing flow. This reduces the label-space shortage by sharing labels for different flows with the same destination, or the same FEC (Forwarding Equivalence Class).

Use the `no` parameter to revert to the default merge capability settings of all the interfaces for this LSR.

Command Syntax

```
global-merge-capability (merge-capable|non-merge-capable)
no global-merge-capability
```

Parameters

merge-capable

Maps all incoming labels that are destined for the same FEC to the same outgoing label (this is the Ethernet default.)

non-merge-capable

Maps all incoming labels, regardless of destination FEC to unique outgoing labels (this is the non-Ethernet default.)

Default

By default, global merge capability is merge capable.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ldp
(config-router)#global-merge-capability merge-capable
```

graceful-restart

Use this command to enable the Graceful-Restart capability for LDP.

Use the `no` parameter to disable the GR capability for LDP.

Command Syntax

```
graceful-restart full
graceful-restart helper-only
graceful-restart timers max-recovery <15-600>
graceful-restart timers neighbor-liveness <5-300>
no graceful-restart
no graceful-restart timers max-recovery
no graceful-restart timers neighbor-liveness
```

Parameters

full

Configuring with full enable the complete GR capability

helper-only

Configuring with helper-only enables only helper mode

timers

Non-default recovery and reconnect timer values.

max-recovery

Maximum recovery time

<15-600>

Interval until which LDP preserves route after peer restart

neighbor-liveness

Neighbor Liveness Time

<5-300>

Hold timer for a targeted LDP peer

Default

GR capability is not enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 5.0.

Examples

```
#configure terminal
(config)#router ldp
(config-router)#graceful-restart full
(config-router)#graceful-restart helper-only
(config-router)#graceful-restart timers max-recovery 100
```

```
(config-router)#graceful-restart timers neighbor-liveness 200
```

hello-interval

Use this command to set the interval after which `hello` packets are sent out.

LDP defines a mechanism for discovering adjacent Label Switching Routers (LSRs) that participate in label switching (adjacencies). Hello messages are sent to the All Routers Multicast Group (224.0.0.2). Whenever a new router comes up, it sends out a hello packet to a specified, multicast address announcing itself to the network. Every router directly connected to the network receives the packet. Receipt of a hello packet from another LSR creates a `hello adjacency` with that LSR. Use this command to specify the interval after which the hello packets will be sent.

Used as a global command, the `hello-interval` value may be overridden by the `hello-interval` set on the interface (see [ldp hello-interval \(page 557\)](#)). For optimum performance, set this value to no more than one-third the value of the hold-time specified.

Use the `no` parameter to revert to default hello interval.

Command Syntax

```
hello-interval <1-21845>
no hello-interval
```

Parameters

<1-21845>

Specify the interval in seconds. The default is 5 seconds.

Default

By default, hello interval is 5 seconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example shows how to set the `hello-interval` value for all interfaces of an LSR.

```
#configure terminal
(config)#router ldp
(config-router)#hello-interval 35

(config-router)#no hello-interval
```

hold-time

Use this command to set the global value for the hold-time after which the LSR rejects adjacencies.

An LSR maintains a record of `hello`s received from peers. `Hold-time` specifies the time an LSR maintains its record of hellos from a peer on not receiving another hello from that peer. A pair of LSRs negotiates the hold-time they use for hellos from each other. Each proposes a hold time value, and the LSR uses the lower of the two hold-time values. The hold-time value set on the interface overrides the hold-time value set by this command (see `ldp hold-time`). For optimum performance, set this value to no less than three times the value of the hello-interval specified.

Use the `no` parameter to revert to the default hold time.

Command Syntax

```
hold-time <3-65535>
no hold-time
```

Parameters

<3-65535>

Specify the hold-time value in seconds.

Default

By default, hold time is 15 seconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to set the hold-time value for all interfaces of an LSR.

```
#configure terminal
(config)#router ldp
(config-router)#hold-time 635

(config-router)#no hold-time
```

import-bgp-routes

Use this command to import BGP routes into LDP. BGP routes are not imported into LDP by default.

Use the `no` parameter to flush out all BGP routes currently being used by LDP, and to reject any further BGP specific routing updates from OcNOS.

Command Syntax

```
import-bgp-routes
no import-bgp-routes
```

Parameters

None

Default

By default, import bgp route is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#import-bgp-routes
```

inter-area-lsp

Use this command to enable creation of inter-area LSPs.

Use the `no` form of the command to disable this configuration.

Command Syntax

```
inter-area-lsp (PREFIX_ACL|) (config-only|)
no inter-area-lsp
```

Parameters

PREFIX_ACL

Access-list name for Prefix Based inter-area lsp

config-only

Optional. When this option is used, existing LDP sessions are not torn down.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 4.0.

Example

```
#configure terminal
(config)#router ldp
(config-router)#inter-area-lsp

#configure terminal
(config)#router ldp
(config-router)#inter-area-lsp config-only

#configure terminal
(config)#router ldp
(config-router)#inter-area-lsp acl1

#configure terminal
(config)#router ldp
(config-router)#inter-area-lsp acl1 config-only
```

keepalive-interval

Use this command to set the global value for the interval after which keep-alive packets are sent out.

Each LSR must send keep-alive messages at regular intervals to its LDP peers to keep the sessions active. The keep-alive interval determines the time interval between successive keep-alive messages. Use this command to set this interval. This value is overridden by the keep-alive interval set on the interface. For optimum performance, set this value to no more than one-third the value of the specified keep-alive time-out value.

Use the `no` parameter to revert to default keep-alive interval.

Command Syntax

```
keepalive-interval <10-21845>  
no keepalive-interval
```

Parameters

<10-21845>

Specify the value of interval in seconds.

Default

By default, keepalive interval is 10 seconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to set the keep-alive timer for all interfaces of an LSR.

```
#configure terminal  
(config)#router ldp  
(config-router)#keepalive-interval 635  
  
(config-router)#no keepalive-interval
```

keepalive-timeout

Use this command to set the global value for the time-out after which sessions are rejected.

Use this command to set the time period for which an LSR must wait for successive keep-alive messages from LDP peers. The keep-alive time-out value is overridden by the keep-alive time-out set on the interface (see `ldp keepalive-timeout`). For optimum performance, set this value to no less than three times the value of the specified keep-alive interval value.

Use the `no` parameter to revert to default keep-alive time-out.

Command Syntax

```
keepalive-timeout <30-65535>  
no keepalive-timeout
```

Parameters

<30-65535>

Specify the time-out value in seconds.

Default

By default, keepalive timeout is 30 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to set the keep-alive time-out value for all interfaces of an LSR.

```
#configure terminal  
(config)#router ldp  
(config-router)#keepalive-timeout 635  
(config-router)#no keepalive-timeout
```

label-retention-mode

Use this command to set the retention mode to be used for all labels exchanged.

When an LSR receives a label binding for a particular FEC (Forwarding Equivalence Class) from another LSR that is not its next hop for that FEC, it might keep track of such bindings or discard them. Use the `liberal` parameter to retain all labels binding to FEC received from label distribution peers, even if the LSR is not the current next-hop. Use the `conservative` parameter to maintain only the label bindings for valid next-hops in a LSP. Liberal label retention mode allows for quicker adaptation to routing changes, whereas conservative label retention mode requires an LSR to maintain fewer labels.



Note: The retention mode value set on the interface (see [ldp label-retention-mode \(page 561\)](#)) overrides the value set by this command.



Note: Any changes made to the retention mode for an interface (after a session is already operational) will only apply to labels received after the mode has been changed. All previously received labels will remain as they were.

Use the `no` parameter to revert to default retention mode.



Note: label-retention-mode "liberal" is supported with advertisement mode "DU" only.



Note: label-retention-mode "liberal" is not supported with advertisement mode "DOD".



Note: When the advertisement mode is set as "DU", label-retention-mode automatically sets to "liberal".



Note: label-retention-mode "conservative" is supported with advertisement mode "DOD" only.



Note: label-retention-mode "conservative" is not supported with advertisement mode "DU".



Note: When the advertisement mode is set as "DOD", label-retention-mode automatically sets to "conservative".

Command Syntax

```
label-retention-mode (conservative|liberal)
no label-retention-mode (conservative|liberal)
```

Parameters

conservative

Specify to delete all unused labels and FECs.

liberal

Specify to retain all labels, regardless of use.

Default

By default, label retention mode is liberal

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to set the retention mode for all interfaces of an LSR.

```
#configure terminal
(config)#router ldp
(config-router)#label-retention-mode liberal
```

ldp advertisement-mode

Use this command to set the label advertisement mode for an interface for the current LSR to either downstream-on-demand (label is sent only when requested) or downstream-unsolicited (label is sent unrequested). Specifying downstream-on-demand and downstream-unsolicited mode affects which LSR initiates mapping requests and mapping advertisements.

This is an interface-specific command; it overrides the advertisement mode set for an LSR using the advertisement-mode command (see [advertisement-mode \(page 522\)](#)). Use this command after the advertisement-mode command sets all the interface advertisement modes. In addition, users should use this command before starting the interface, since all affected sessions will be closed and restarted.

Use the `no` parameter to revert to the advertisement mode value set for the main LDP process.

Command Syntax

```
ldp advertisement-mode (downstream-on-demand|downstream-unsolicited)
no ldp advertisement-mode (downstream-on-demand|downstream-unsolicited)
```

Parameters

downstream-on-demand

Indicates that the sent label was requested. When a user uses this parameter, a router distributes a label to a peer only if there is a pending label request from a peer. The reaction of the downstream router to this request depends on the label advertising mode supported on the next hop. The downstream-on-demand mode is typically used with the conservative label retention mode.

downstream-unsolicited

Indicates that the label was sent unrequested. This parameter distributes labels to peers without waiting for a label request. This mode is typically used with the liberal label retention mode.

Default

By default, ldp advertisement mode is downstream unsolicited mode

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#ldp advertisement-mode downstream-on-demand
```


ldp hello-interval

Use this command to set the interval for sending multicast Hello packets via an interface.

LDP defines a mechanism for discovering adjacent Label Switching Routers (LSR) that participate in label switching (adjacencies). Whenever a new router comes up, it sends out a hello packet to a specified, multicast address announcing itself to the network. Every router directly connected to the network receives the packet. Receipt of a hello packet from another LSR creates a hello adjacency with that LSR. Use this command to specify the interval after which the hello packets will be sent.

For optimum performance, set the hello-interval value to no more than one-third the hold-time value.



Note: This command is an interface-specific command and overrides the value set for an LSR using the global hello-interval command.

Use the `no` parameter with this command to revert to the hello-interval value set for the main LDP process.

Command Syntax

```
ldp hello-interval <1-21845>
no ldp hello-interval
```

Parameters

<1-21845>

Specify the interval in seconds.

Default

By default, ldp hello interval is 5 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example shows how to set the hello-interval for a specific interface.

```
#configure terminal
(config)#interface eth0
(config-if)#ldp hello-interval 635
(config-if)#no ldp hello-interval
```

ldp hold-time

Use this command to set the hold-time value after which the LSR rejects adjacencies.

The hold-time timer is reset every time a hello packet is received from the peer in question. For optimum performance, set this value to no less than three times the hello-interval value.



Note: This command is an interface-specific command, and overrides the value set for an LSR using the global hold-time command.

Use the `no` parameter to revert to the hold-time value set for the main LDP process.

Command Syntax

```
ldp hold-time <3-65535>
no ldp hold-time
```

Parameters

<3-65535>

Specify the hold-time value in seconds.

Default

By default, ldp hold time is 15 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to set the hold-time for a specific interface:

```
#configure terminal
(config)#interface eth0
(config-if)#ldp hold-time 635
(config-if)#no ldp hold-time
```

ldp keepalive-interval

Use this command to set the interval for sending keep-alive messages to the peer in order to maintain a session. Each LSR must send keep-alive messages at regular intervals to its LDP peers to keep the sessions active. The keep-alive interval determines the time-interval between successive keep-alive messages. This command sets this interval.



Note: This command is an interface-specific command, and overrides the value set for an LSR using the global `keepalive-interval` command.

Use the `no` parameter to revert to the keep-alive interval set for the main LDP process.

Command Syntax

```
ldp keepalive-interval <1-21845>
no ldp keepalive-interval
```

Parameters

<1-21845>

Specify the interval in seconds.

Default

By default, ldp keepalive interval is 10 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

This example shows how to set the hello-interval for a specific interface:

```
#configure terminal
(config)#interface eth0
(config-if)#ldp keepalive-interval 635
(config-if)#no ldp keepalive-interval
```

ldp keepalive-timeout

Use this command to set the keep-alive time-out value for rejecting a session with a peer.

Use this command to set the time period for which an LSR must wait for successive keep-alive messages from LDP peers. The keep-alive timer is reset every time a keep-alive packet is received from the peer in question. For optimum performance, set this value to no more than three times the keep-alive interval value.



Note: This command is an interface-specific command and overrides the value set for an LSR using the global `keepalive-timeout` command.

Use the `no` parameter to revert to the keep-alive time-out set for the main LDP process.

Command Syntax

```
ldp keepalive-timeout <3-65535>  
no ldp keepalive-timeout
```

Parameters

<3-65535>

Specify the value in seconds.

Default

By default, ldp keepalive timeout is 30 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to set the keep-alive time-out timer for a specific interface:

```
#configure terminal  
(config)#interface eth0  
(config-if)#ldp keepalive-timeout 635  
(config-if)#no ldp keepalive-timeout
```

ldp label-retention-mode

Use this command to set the retention mode to be used for all labels exchanged via the given interface.

When an LSR receives a label binding for a particular FEC (Forwarding Equivalence Class) from another LSR that is not its next hop for that FEC, it might keep track of such bindings or discard them. Use the `liberal` parameter to retain all labels binding to FEC received from label distribution peers, even if the LSR is not the current next-hop. Use the `conservative` parameter to maintain only the label bindings for valid next-hops in a LSP. Liberal label retention mode allows for quicker adaptation to routing changes, whereas conservative label retention mode requires an LSR to maintain fewer labels.



Note: The retention mode value set on the interface (see [label-retention-mode \(page 554\)](#)) overrides the value set by this command. This command is an interface-specific command, and overrides the setting for an LSR using the global `label-retention-mode` command.

Use the `no` parameter to revert to the retention mode set for the main LDP process.

Command Syntax

```
ldp label-retention-mode (conservative|liberal)
no ldp label-retention-mode (conservative|liberal)
```

Parameters

conservative

Specify to delete all unused labels and FECs.

liberal

Specify to retain all labels, regardless of use.

Default

By default, ldp label retention mode is liberal

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows how to set the label retention mode for a specific interface:

```
#configure terminal
(config)#interface eth0
(config-if)#ldp label-retention-mode liberal
```

ldp multicast-hellos

Use this command to enable multicast hello exchange on a specified interface.

Use the `no` parameter to disable multicast hello exchange. R

Command Syntax

```
ldp multicast-hellos  
no ldp multicast-hellos
```

Parameters

None

Default

By default, ldp multicast hello is enabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#interface eth0  
(config-if)#ldp multicast-hellos
```

ldp-optimization

This command helps optimize the resetting of an LDP session by enabling the following two scalability features for LDP:

- Resets the session keepalive timer on receipt of a hello message
- Resets the hold timer on receipt of any LDP control message

Use the `no` parameter to disable the two previously listed scalability features.

Command Syntax

```
ldp-optimization
no ldp-optimization
```

Parameters

None

Default

By default, ldp optimization is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#ldp-optimization
```

loop-detection

Use this command to enable loop detection on the current LSR. This command detects looping LSPs, and prevent Label Request messages from looping because of non-merge capable LSRs. This loop detection mechanism is useful for networks of non time-to-live (non TTL) decrementing devices that can not allocate resources among traffic flows.

There are two methods supported for the loop detection mechanism: A Hop Count detection system, that is always enabled; and the Path Vector detection system, that can be toggled:

- Hop Count - During the setup of an LSP, the LSP passes a hop count with the LSP setup messages. This hop count is incremented by each node router participating in LSP establishment. If the hop count exceeds the maximum configured value, the LSP setup process is stopped, and a notification message is passed back to the message originator.
- Path Vector - A path vector contains a list of LSR identifiers. This is passed as a part of LSP setup messages. Each LSR participating in the LSP establishment adds its own LSR identifier to the path vector. If an LSR finds its own identifier in the path vector, it drops the message, and sends a message back to the originator.

The use of these messages ensures that a loop is detected while establishing a label switched path and before any data is passed over that LSP.

Use the `no` parameter to disable loop detection.

Command Syntax

```
loop-detection
no loop-detection
```

Parameters

None

Default

By default, loop detection is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#loop-detection
```

loop-detection-hop-count

Use this command to set the loop detection hop count, which determines the maximum hop-count value.

This command sets the maximum hop count value, which specifies the permitted maximum permitted hop-count. An LSR that detects a maximum hop count behaves as if the containing message has traversed a loop. The use of this command ensures that a loop is detected while establishing a label switched path before any data is passed via LSP.

Use the `no` parameter to revert to the default loop detection count

Command Syntax

```
loop-detection-hop-count <1-255>
```

Parameters

<1-255>

Indicates the loop detection hop count.

Default

By default, loop detection hop is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ldp
(config-router)#loop-detection-hop-count 128
```

loop-detection-path-vec-count

Use this command to set the loop detection vec (vector) count, which determines the maximum supported path vectors.

This command sets the maximum supported path vectors for loop detection, which specifies the permitted path vector length. An LSR that detects a path vector has reached the maximum length behaves as if the containing message has traversed a loop. This command ensures that a loop is detected while establishing a label switched path before any data is passed over that LSP.

Use the `no` parameter to revert to the default loop detection count

Command Syntax

```
loop-detection-path-vec-count <1-255>
```

Parameters

<1-255>

Indicates the loop detection hop count.

Default

By default, loop detection path vec count is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#loop-detection-path-vec-count 123
```

mpls ldp-igp sync-delay

Use this command to set the time delay for LDP-IGP synchronization.

Use the `no` parameter to disable the time delay.

Command Syntax

```
mpls ldp-igp sync-delay <5-60>
no mpls ldp-igp sync-delay
```

Parameters

sync-delay

Time delay for LDP to converge in seconds.

<5-60>

Time delay for notification of LDP convergence to IGP, in seconds

Default

If not configured the delay will be 0 seconds.

Command Mode

Interface configuration mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config-if)# interface eth0
(config-if)# mpls ldp-igp sync-delay 15
(config-if)# no mpls ldp-igp sync-delay
```

mpls ldp-igp sync isis

Use this command to enable LDP ISIS synchronization and to set the holddown timer for synchronization.

Use the `no` parameter to disable the LDP ISIS synchronization.



Note: Holddown timer value should be higher than LDP IGP sync timer.

Command Syntax

```
mpls ldp-igp sync isis (level-1|level-2|level-1-2) (holddown-timer <1-2147483>| )
```

Parameters

level-1|level-2|level-1-2

The ISIS level.

holddown-timer

How long IGP should wait for LDP to converge in seconds.

Default

None

Command Mode

Interface configuration mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
#int eth 1
#mpls ldp-igp sync isis level-1-2 holddown-timer 500
```

mpls ldp-igp sync ospf

Use this command to enable LDP-OSPF synchronization. This command also provides option to configure the hold-down timer for which OSPF will wait for LDP to converge and advertises Max cost. When the configured time expires, OSPF starts advertising the actual cost in the Router-LSA.



Note: Holddown timer value should be higher than LDP IGP sync timer.

Command Syntax

```
mpls ldp-igp sync ospf (holddown-timer <1-2147483>|)
```

Parameters

holddown-timer

Set holddown timer for the OSPF Sync

<1-2147483>

Hold down timer in seconds

Default

OSPF waits infinite when no hold-down timer is configured.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 4.0.

Example

Enabling OSPF-LDP sync in interface eth3

```
#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#int xe3
(config-if)#mpls ldp-igp sync ospf
(config-if)#end
```

Enabling OSPF-LDP sync with holddown-timer enabled

```
#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#int xe3
(config-if)#mpls ldp-igp sync ospf holddown-timer 200
(config-if)#no mpls ldp-igp sync ospf
(config-if)#end
#
```

multicast-hellos

Use this command to enable multicast hello exchange on all interfaces enabled for LDP. This is used for auto-discovery of LDP peers on directly connected networks. This option is enabled by default.

Use the `no` parameter with this command to disable multicast hello exchange.

Command Syntax

```
multicast-hellos  
no multicast-hellos
```

Parameters

None

Default

By default, multicast hello is enabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#router ldp  
(config-router)#multicast-hellos
```

neighbor

Use this command to configure neighbors of LDP.

Use the `no` parameter with this command to unconfigure the LDP neighbor.

Command Syntax

```
neighbor A.B.C.D auth AUTH-TYPE password (0|7) WORD
no neighbor A.B.C.D auth AUTH-TYPE password
```

Parameters

A.B.C.D

Neighbor address

auth AUTH-TYPE

Authentication Type md5

password

Set password to the neighbor

(0|7)

Password Type

WORD

Password

Default

By default, neighbor is disabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#neighbor 1.1.1.1 auth md5 password 0 myPass

(config-router)#no neighbor 1.1.1.1 auth md5 password
```

neighbor auth md5 password

Use this command to configure LDP neighbors for MD5 authentication.

Use the `no` parameter with this command to unconfigure the LDP neighbors for MD5 authentication.

Command Syntax

```
neighbor (all|auto-targeted|A.B.C.D) auth md5 password (plain-text|encrypt) WORD
no neighbor (all|auto-targeted|A.B.C.D) auth md5 password

neighbor (A.B.C.D|all|auto-targeted) auth md5 password exclude
no neighbor (A.B.C.D|all|auto-targeted) auth md5 password exclude
```

Parameters

all

Enable MD5 authentication for all LDP neighbors.

auto-targeted

Enable MD5 authentication to auto-targeted neighbors only.

A.B.C.D

Enable MD5 authentication to a specific neighbor using its IPv4 address.

password plain-text

Enter the password in plain text. This option stores the password as clear text, making it easily readable.

password encrypt

Use an encrypted password. This option encrypts the password while storing it, preventing unauthorized access.

WORD

Specify the actual password (either in plain text or encrypted).

password exclude

Set the neighbor to exclude password authentication.

Default

None

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#auto-targeted-session
(config-router)#neighbor auto-targeted auth md5 password plain-text test2
```


neighbor prefix-list

Use this command to configure the peer authentication group prefix list.

Command Syntax

```
neighbor prefix-list NAME
```

Parameters

NAME

Specify the peer authentication group prefix list name.

Default

None

Command Mode

Session group mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#session-group name abcd
(config-router-sg)#neighbor prefix-list test
```

propagate-release

Use this command to propagate the release of labels to downstream routers.

Use the `no` parameter to prevent the propagate-release of labels.

Command Syntax

```
propagate-release  
no propagate-release
```

Parameters

None

Default

By default, propagate release is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#router ldp  
(config-router)#propagate-release
```

pw-status-tlv

Use this command to enable the use of the PW Status TLV to signal the pseudowire status.

Use the no option with this command to disable the use of the PW Status TLV to signal the pseudowire status.

Command Syntax

```
pw-status-tlv  
no pw-status-tlv
```

Parameters

None

Default

By default, pw status tlv is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#router ldp  
(config-router)#pw-status-tlv
```

request-labels-for

Use this command to request labels for the prefixes in the given IP prefix list. LDP request labels for the prefixes only if the valid and exact route is present for that prefix.

Use the no form of this command to disable multicast hello exchange.

Command Syntax

```
request-labels-for prefix-list-ipv4 NAME
no request-labels-for prefix-list-ipv4
```

Parameters

NAME

IPv4 prefix list name

Command Mode

LDP router mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#router ldp
(config-router)#request-labels-for prefix-list-ipv4 myPrefixList
```

request-retry

Use this command to enable the retry of requests once a request for a label has been rejected for a valid reason. This command enables the LSR to send a maximum of five label requests if a label request is rejected by an LDP peer.

Use the `no` parameter to disable the retry of requests.

Command Syntax

```
request-retry
no request-retry
```

Parameters

None

Default

By default, request retry is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router ldp
(config-router)#request-retry
```

request-retry-timeout

Use this command to set the interval between retries. Before this time is over, a request is re-sent to a peer. This command changes the interval between request messages that are resent to a peer to account for routing changes. Use the `no` parameter to revert to the default request-retry time-out set.

Command Syntax

```
request-retry-timeout <1-65535>
no request-retry-timeout
```

Parameter

<1-65535>

Specify the interval between retries in seconds.

Default

By default, timeout is 5 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#request-retry-timeout 512

(config-router)#no request-retry-timeout
```

restart ldp graceful

Use this command to restart ldp gracefully.

Command Syntax

```
restart ldp graceful
```

Parameter

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 5.0.

Example

```
#restart ldp graceful

% Warning : LDP process will stop and needs to restart manually,
You may loose LDP configuration, if not saved
Proceed for graceful restart? (y/n):y
```

rlfa-ilm-optimization

Use this command to enable ILM optimization for rLFA. When ILM optimization is enabled, ILM delete and POP operations are delayed. ILM delete/POP configuration is delayed to make sure rLFA backup path is not removed immediately upon IGP network convergence.

Use the `no` parameter with this command to revert this configuration.

Command Syntax

```
rlfa-ilm-optimization
no rlfa-ilm-optimization
```

Parameter

None

Default

By default, ILM optimization is disabled.

Command Mode

Router(LDP) Config mode.

Applicability

This command was introduced before OcNOS version 6.2.0.

Example

```
(config)#router ldp
(config-router)# rlfa-ilm-optimization
```

router-id

Use this command to set the router-id to the supplied IP address; the router uses this address to generate the LDP-ID.

OcNOS has three methods to choose the router-id of LDP. The first priority router-id is the configured router-id in router mode (local configured router-id). The second priority router-id is the configured router-id in configure mode (global configured router-id). The lowest priority router-id is chosen by OcNOS among interfaces (global computed router-id).

Use the `no` parameter with this command to revert to using the first IP address configured on the box as the router-id for LDP-ID generation purposes.

Command Syntax

```
router-id A.B.C.D
no router-id A.B.C.D
no router-id
```

Parameter

A.B.C.D

Indicates the LDP router ID value.

Default

By default, router id is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure router
(config)#router ldp
(config-router)#router-id 123.123.123.8
```

router ldp

This command is used to enter the LDP specific command-line mode in which global attributes for the LDP process can be set. Without this command, the LSR does not perform any LDP operations, such as sending `hello` packets. Use the `no` parameter with this command to disable this configuration.

Command Syntax

```
router ldp
no router ldp
```

Parameters

None

Default

By default, router ldp is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following example shows the change in the prompt after using this `router ldp` command to enter router mode.

```
#configure router
(config)#router ldp
(config-router)#
```

session-group name

Use this command to create an LDP session group name.

Command Syntax

```
session-group name NAME
```

Parameters

Name

Specify the LDP session group name.

Default

None

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#session-group name test
```

snmp restart ldp

Use this command to restart SNMP in Label Distribution Protocol (LDP)

Command Syntax

```
snmp restart ldp
```

Parameters

None

Default

By default, snmp restart ldp is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#snmp restart ldp
```

targeted-peer hello-interval

Use this command to set the hold time for a targeted LDP peer.

Use the **no** parameter with this command to unset hold time for a targeted LDP peer.

Command Syntax

```
hello-interval <1-21845>  
no hello-interval
```

Parameters

<3-65535>

Hold time in seconds as time-out value that the router waits before rejecting an adjacency.

Default

45

Command Mode

Targeted-peer mode

Applicability

This command was introduced before OcNOS version 6.5.3.

Examples

This example shows how to set the hello-interval for a specific interface.

```
ocnos(config)#router ldp  
ocnos(config-router)#targeted-peer ipv4 1.1.1.1  
ocnos(config-router-targeted-peer)# hello-interval 10
```

targeted-peer ipv4

Use this command to enter a targeted IPv4 LDP peer mode.

A targeted session is an LDP session between non-directly connected LSRs. Set this command to send a targeted hello messages to specific IP addresses. This command is specific to a targeted IPv4 LDP peer.

Command Syntax

```
targeted-peer ipv4 A.B.C.D  
no targeted-peer ipv4 A.B.C.D
```

Parameter

A.B.C.D

Specify the IPv4 address of the targeted peer.

Default

By default, targeted peer IPv4 is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#router ldp  
(config-router)#targeted-peer ipv4 10.10.10.10  
(config-router-targeted-peer)#
```

transport-address ipv4

Use this command to configure the IPv4 transport address for a label space.

The transport address is the address used for the TCP session over which LDP is running. Use this command to manually configure the transport address. Transport addresses may either be bound to a loopback interface, or to a physical interface that is bound to the label space in question. A transport address can also be manually configured using the CLI with the loopback address as the transport address.



Note: The CLI accepts only the loopback address to be configured as the transport address.

Use the `no` parameter to stop using the transport address as the IPv4 transport address. If the label space is not specified for either form of this command, a label space of zero is assumed.

Command Syntax

```
transport-address ipv4 A.B.C.D
transport-address ipv4 A.B.C.D 0
no transport-address ipv4 A.B.C.D
no transport-address ipv4 A.B.C.D 0
```

Parameters

A.B.C.D

Specify the IPv4 address to be used as the transport address. Only addresses bound to a loopback interface are valid for manual transport address configuration.

0

Platform-wide label space for which a transport address is being configured (Platformwide labels are used for all interfaces that can share the same labels)

Default

Transport addresses are chosen for label spaces. By default, the loopback address is selected as the transport address. If a loopback address is not configured, the label space value is examined. The IP address of the interface is bound to the same label space is chosen as the transport address.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure router
(config)#router ldp
(config-router)#transport-address ipv4 10.10.0.5 0
```

targeted-peer tunneling

Use this command to enable the TLDP peer to tunnel LDP FECs.

Use the **no** parameter with this command to disable the TLDP peer from tunneling LDP FECs.

Command Syntax

```
tunneling
no tunneling
```

Parameters

None

Default

Disable

Command Mode

Targeted-peer mode

Applicability

This command was introduced before OcNOS version 6.6.1.

Examples

This example shows how to tunnel LDP FECs to TLDP peers.

```
Enter configuration commands, one per line. End with CNTL/Z.
ocnos(config)#router ldp
ocnos(config-router)# targeted-peer ipv4 2.2.2.2
ocnos(config-router-targeted-peer)#tunneling
ocnos(config-router-targeted-peer)#commit
```

targeted-peer hold-time

Use this command to set the hold time for a targeted LDP peer.

Use the **no** parameter with this command to unset hold time for a targeted LDP peer.

Command Syntax

```
hold-time <3-65535>  
no hold-time
```

Parameters

<1-21845>

Interval for sending periodic hello message to peers.

Default

15

Command Mode

Targeted-peer mode

Applicability

This command was introduced before OcNOS version 6.5.3.

Examples

This example shows how to set the hello-interval for a specific interface.

```
ocnos(config)#router ldp  
ocnos(config-router)#targeted-peer ipv4 1.1.1.1  
ocnos(config-router-targeted-peer)# hold-time 4
```

targeted-peer-hold-time

Use this command to set the time-out value that is the time that the router waits before rejecting an adjacency with targeted peers.

Use the `no` parameter to revert to the default targeted-peer hold-time value.

Command Syntax

```
targeted-peer-hold-time <3-65535>
no targeted-peer-hold-time
```

Parameter

<3-65535>

Specify the interval in seconds.

Default

By default, hold time is 45 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router ldp
(config-router)#targeted-peer-hold-time 555

(config-router)#no targeted-peer-hold-time
```

LDP Show Commands

This chapter provides an alphabetized reference for each of the LDP commands. It includes the following commands:

| | |
|--|-----|
| show ldp label-pool | 592 |
| show debugging ldp | 593 |
| show ldp | 594 |
| show ldp adjacency | 596 |
| show ldp adjacency count | 597 |
| show ldp advertise-labels | 598 |
| show ldp fec | 599 |
| show ldp igp sync | 601 |
| show ldp inter-area-fecs | 602 |
| show ldp inter-area-fecs prefix | 603 |
| show ldp interface | 604 |
| show ldp lsp | 606 |
| show ldp mpls-l2-circuit | 609 |
| show ldp rfa-routes | 612 |
| show ldp rfa-routes count | 613 |
| show ldp routes | 614 |
| show ldp routes count | 616 |
| show ldp session | 617 |
| show ldp session all count | 620 |
| show ldp session multicast count | 621 |
| show ldp session targeted count | 622 |
| show ldp statistics | 623 |
| show ldp statistics advertise-labels | 625 |
| show ldp targeted-peers | 626 |
| show ldp upstream | 627 |
| show ldp downstream | 629 |
| show ldp vpls | 631 |
| show mpls ldp discovery | 634 |
| show mpls ldp neighbor | 635 |
| show mpls ldp parameter | 636 |
| show ldp tunneling-fec | 638 |
| show ldp tunneling | 639 |
| show ldp tunneling-tunnels | 640 |

show ldp label-pool

Use this command to display the label management in LDP protocol.

Command Syntax

```
show ldp label-pool (block-id <0-1638>|)
```

Parameter

block-id

Block identifier of a block that is allocated to the protocol.

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 6.5.4

Example

The following is an output from the show ldp label-pool command.

```
show ldp label-pool
Module: LDP, Label range: 16-1048575, Current block: 40
+-----+-----+-----+-----+-----+-----+
| block_id | label_min | label_max | usable_labels | free_labels | first_free_label |
+-----+-----+-----+-----+-----+-----+
| 40 | 25600 | 26239 | 640 | 628 | 25612 |
Total - blocks: 1, used-labels: 12, free-labels: 628
Module: LDP-VC, Label range: 16-1048575, Current block: 79
+-----+-----+-----+-----+-----+-----+
| block_id | label_min | label_max | usable_labels | free_labels | first_free_label |
+-----+-----+-----+-----+-----+-----+
| 41 | 26240 | 26879 | 640 | 0 | 0 |
| 79 | 50560 | 51199 | 640 | 456 | 50744 |
Total - blocks: 2, used-labels: 824, free-labels: 456
```

show debugging ldp

Use this command to display the status of the debugging of the LDP system.

Command Syntax

```
show debugging ldp
```

Parameter

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output from the `show debugging ldp` command.

```
#show debugging ldp
LDP debugging status:
  LDP event debugging is on
  LDP packet debugging is on
  LDP finite state machine debugging is on
  LDP pdu hexdump debugging is on
  LDP downstream state machine debugging is on
  LDP upstream state machine debugging is on
  LDP trunk state machine debugging is on
  LDP QoS debugging is on
  LDP CSPF debugging is on
  LDP VC USM debugging is on
  LDP VC DSM debugging is on
  LDP NSM debugging is on
  LDP Advertise-labels debugging is on
#
```

[Table 8](#) explains the show command output fields.

Table 8. show debugging ldp output fields details

| Field | Description |
|----------------------|---------------------------------------|
| LDP debugging status | Status of the LDP debugging protocol. |

show ldp

Use this command to display basic LDP attributes defined for the current LSR.

Command Syntax

```
show ldp
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following is a sample output from the `show ldp` command displaying basic LDP attributes.

```
#show ldp
Router ID : 20.1.1.1
LDP Version : 1
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode : Liberal
Label Control Mode : Independent
Instance Loop Detection : On
Instance Hop Count Limit : 255
Instance Path Vec Count : 255
Request Retry : Off
Propagate Release : Disabled
Graceful Restart : Disabled
Hello Interval : 5
Targeted Hello Interval : 15
Hold time : 15
Targeted Hold time : 45
Keepalive Interval : 10
Keepalive Timeout : 30
Request retry Timeout : 5
Transport Address data :
Labelspace 0 : 20.1.1.1 (in use)
Import BGP routes : No
Entropy-Label Capability : Disabled
Prefer-Tunnel-In-Tunnel : Disabled
#
```

[Table 9](#) explains the show command output fields.

Table 9. show ldp output fields details:

| Field | Description |
|--------------------------|---|
| Router ID | Router identifier in IP address format for this system. |
| LDP Version | Details of Link Layer Discovery Protocol (LLDP) version. |
| Global Merge Capability | Used to override the default merge capability setting of all the interfaces for the current LSR. |
| Label Advertisement Mode | Used to set the label advertisement mode for an interface for the current LSR to either downstream-on-demand (label is sent only when requested) or downstream-unsolicited (label is sent unrequested). |
| Label Retention Mode | Used for all labels exchanged via the given interface. |
| Label Control Mode | LSR generates a local label for a FEC which the router learned from routing table independently from other LSRs. |
| Loop Detection | Used to enable loop detection on the current LSR. |
| Loop Detection Count | Indicates the loop detection hop count. |
| Request Retry | Enables the LSR to send a maximum of five label requests. |
| Propagate Release | Used to propagate the release of labels to downstream routers. |
| Hello Interval | Sets the interval for sending unicast hello packets to peers. |
| Targeted Hello Interval | Sets the interval for sending unicast hello packets to targeted peers. |
| Hold time | Sets the time-out value to peers. |
| Targeted Hold time | Sets the time-out value that is the time that the router waits before rejecting an adjacency with targeted peers. |
| Keepalive Interval | Used to set the interval for sending keep-alive messages to the peer in order to maintain a session. |
| Keepalive Timeout | Time-out value for rejecting a session with a peer. |
| Request retry Timeout | Used to set the interval between retries. |
| Targeted Hello Receipt | Status of the hello receipt. |
| Transport Address | The transport address is the address used for the TCP session over which LDP is running. |
| Transport Interface | Interface is used for the TCP session over which LDP is running. |
| Import BGP routes | Used to import BGP routes into LDP. |
| Entropy-Label Capability | Displays the status of the Entropy-Label Capability capability. The status can be either enabled or disabled. |
| Prefer-Tunnel-In-Tunnel | Displays the status of the Prefer-Tunnel-In-Tunnel capability. The status can be either enabled or disabled for LDP over RSVP. |

show ldp adjacency

Use this command to display all the adjacencies for the current LSR.

Command Syntax

```
show ldp adjacency
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output from the `show ldp adjacency` command displaying all the adjacencies for this LSR.

```
#show ldp adjacency
Remote-Address  Local-Address  Mode          Intf-Name  Holdtime  LDP-Identifier
11.11.11.11     12.0.1.20     Targeted      ge11       45        11.11.11.11:0
33.33.33.33     11.0.1.20     Targeted      ge9        45        33.33.33.33:0
44.44.44.44     20.0.1.20     Targeted      xe14       45        44.44.44.44:0
11.0.1.10       11.0.1.20     Interface     ge9        15        33.33.33.33:0
12.0.1.10       12.0.1.20     Interface     ge11       15        11.11.11.11:0
20.0.1.10       20.0.1.20     Interface     xe14       15        44.44.44.44:0
```

[Table 10](#) explains the show command output fields.

Table 10. show ldp adjacency output fields details

| Field | Description |
|----------------|-------------------------------------|
| Remote Address | IP address of the interface. |
| Local Address | Local address of the LDP adjacency. |
| Interface Name | Name of the interface. |
| Hold time | Sets the time-out value to peers. |
| LDP ID | LDP identifier for this protocol. |

show ldp adjacency count

Use this command to display the number of LDP entities, which are enabled interfaces for multicast or targeted peers, and the number of LDP adjacencies created for these peers.

Command Syntax

```
show ldp adjacency count
```

Parameter

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

This command displays the number of up and down multicast LDP sessions

```
OcNOS#show ldp adjacency count

Entity - LDP enabled interfaces / configured targeted-peer
Adjacency -- Created adjacency, once hello received

Type          Entity    Adjacency
-----
Multicast     3          2
Targeted      0          0
Total         3          2
-----
```

[Table 11](#) explains the show command output fields.

Table 11. show ldp session output fields details

| Field | Description |
|--------------------|--|
| session up count | Displays the number of running LDP sessions. |
| session down count | Displays the number of failed LDP sessions. |
| total count | Displays the total number of LDP sessions. |

show ldp advertise-labels

Use this command to display the IP access list of LDP advertise-labels.

Command Syntax

```
show ldp advertise-labels
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output from the `show ldp advertise-labels` command.

```
#show ldp advertise-labels
Advertisement spec:
  Prefix list = pfx1; Peer plist = pfx1
  Prevent the distribution of any assigned labels
```

[Table 12](#) explains the show command output fields.

Table 12. show ldp advertise-labels output fields details

| Field | Description |
|--------------------|--|
| Advertisement spec | Details of the advertisement spec. |
| Prefix list | The label is advertised to all peers permitted by the peer plist. Name of a configured IP prefix-list that specifies the set of prefixes (destinations) to advertise labels for. |
| Peer plist | The prefix list permits the prefix and there is a peer plist. Name of a configured IP prefix-list that specifies the LDP neighbors that receive the label advertisements. |

show ldp fec

Use the following command to display all FECs (Forwarding Equivalence Classes) known to this LSR.

Command Syntax

```
show ldp fec
show ldp fec (prefix)
show mpls ldp fec
show mpls ldp fec (prefix|)
```

Parameter

prefix

Display prefix FEC information.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3 and was updated in OcNOS version 5.1.

Example

When LDP LFA FRR in not enabled:

```
#show ldp fec
LSR codes      : E/N - LSR is egress/non-egress for this FEC,
                  L - LSR received a label for this FEC,
                  > - LSR will use this route for the FEC

FEC             Code   Session      Out Label   ELC   Nexthop Addr
10.0.1.0/31      NL     22.1.1.1    impl-null   No    connected
                  E >   non-existent none        No    connected
10.0.2.0/31      NL     44.1.1.1    24322       No    no nexthop
                  NL>   22.1.1.1    impl-null   No    10.0.1.1
10.0.3.0/31      NL     22.1.1.1    24323       No    no nexthop
                  NL>   44.1.1.1    impl-null   No    10.0.4.1
10.0.4.0/31      NL     44.1.1.1    impl-null   No    connected
                  E >   non-existent none        No    connected
11.1.1.1/32      E >   non-existent none        No    connected
22.1.1.1/32      NL>   22.1.1.1    impl-null   No    10.0.1.1
33.1.1.1/32      NL>   44.1.1.1    24323       No    10.0.4.1
                  NL>   22.1.1.1    24324       No    10.0.1.1
44.1.1.1/32      NL>   44.1.1.1    impl-null   No    10.0.4.1
192.168.254.0/24 NL     22.1.1.1    impl-null   No    connected
                  NL     44.1.1.1    impl-null   No    connected
                  E >   non-existent none        No    connected
```

When LDP LFA FRR in enabled:

```
#show ldp fec
LSR codes      : E/N - LSR is egress/non-egress for this FEC,
                  L - LSR received a label for this FEC,
                  P - Primary route, B - LFA Backup route,
                  R - Remote LFA Backup route,
                  > - LSR will use this route for the FEC

FEC             Code   Session      Out Label   ELC   Nexthop Addr
```

```

10.0.1.0/31      NL      22.1.1.1      impl-null      No      connected
                  E >      non-existent      none           No      connected
10.0.2.0/31      NLB>    44.1.1.1      24322          No      10.0.4.1
                  NLP>    22.1.1.1      impl-null      No      10.0.1.1
10.0.3.0/31      NLB>    22.1.1.1      24323          No      10.0.1.1
                  NLP>    44.1.1.1      impl-null      No      10.0.4.1
10.0.4.0/31      NL      44.1.1.1      impl-null      No      connected
                  E >      non-existent      none           No      connected
11.1.1.1/32      E >      non-existent      none           No      connected
22.1.1.1/32      NLP>    22.1.1.1      impl-null      No      10.0.1.1
33.1.1.1/32      NLP>    44.1.1.1      24323          No      10.0.4.1
                  NLB>    44.1.1.1      24323          No      10.0.4.1
                  NLP>    22.1.1.1      24324          No      10.0.1.1
                  NLB>    22.1.1.1      24324          No      10.0.1.1
44.1.1.1/32      NLP>    44.1.1.1      impl-null      No      10.0.4.1
1.1.1.1/32       NL      1.1.1.1      impl-null      No      no nexthop
                  NLP>    3.3.3.3      24320          No      30.1.1.1
                  NLB>    3.3.3.3      24320          No      30.1.1.1
                  NLP>    2.2.2.2      24320          No      20.1.1.1
                  NLB>    2.2.2.2      24320          No      20.1.1.1
2.2.2.2/32       NLR>    1.1.1.1      24324          No      1.1.1.1
                                (via 30.1.1.1, label 24320)
                  NLP>    2.2.2.2      impl-null      No      20.1.1.1
192.168.254.0/24 NL      22.1.1.1      impl-null      No      connected
                  NL      44.1.1.1      impl-null      No      connected
                  E >      non-existent      none           No      connected

```

[Table 13](#) shows the codes at the end of each route entry that indicate where the route originated.

Table 13. Origin Codes

| Origin Code | Description | Comments |
|-------------|-------------------------|--|
| E/N | Egress/Non-egress | LSR is egress/non-egress for this FEC. |
| L | LSR | LSR received a label for this FEC. |
| > | | LSR will use this route for the FEC. |
| P | Primary route | When LDP LFA FRR is enabled |
| B | LFA Backup route | When LDP LFA FRR is enabled |
| R | Remote LFA Backup route | When LDP LFA FRR is enabled |

[Table 14](#) explains the show command output fields.

Table 14. show ldp fec output fields details

| Field | Description |
|--------------|--|
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| Session | Reports the current session state. |
| Out Label | Label received from downstream neighbor for route. |
| ELC | Displays if route has ELC |
| Nexthop addr | Displays the IP address of the next hop. |

show ldp igp sync

Use the following command to display the LDP synchronization status.

Command Syntax

```
show ldp igp sync
show mpls ldp igp sync
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

| Interface | Link-State | LDP-Enabled | IGP-Sync | Peer-IP | Session-State | Sync-Delay- |
|---------------------------|------------|-------------|----------|---------|---------------|------------------|
| Time/Remaining-Delay-Time | | | | | | |
| xe2 | Up | Yes | Disabled | 2.2.2.2 | Achieved | Not-Configured / |
| Not-Running | | | | | | |
| xe13 | Up | Yes | Disabled | 4.4.4.4 | Achieved | Not-Configured / |
| Not-Running | | | | | | |
| xe16 | Up | Yes | Disabled | 2.2.2.2 | Achieved | Not-Configured / |
| Not-Running | | | | | | |

show ldp inter-area-fecs

Use this command to show all FECs using the LPM-based mapping procedure.

Command Syntax

```
show ldp inter-area-fecs
show ldp inter-area-fecs (ipv4|ipv6|) (count)
```

Parameter

ipv4

IPv4 FECs

ipv6

IPv6 FECs

count

Count of FECs

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 4.0 and the command was updated in OcNOS version 4.1.

Example

```
#show ldp inter-area-fecs
LSR codes : E/N - LSR is egress/non-egress for this FEC,
L - LSR received a label for this FEC,
> - LSR will use this route for the FEC
Code FEC Session Out Label Nexthop Addr
Matching RIB prefix - 1.1.1.0
NL> 1.1.1.1/32 33.33.33.33 52485 11.11.11.1
NL> 1.1.1.2/32 33.33.33.33 52486 11.11.11.1
```

```
#show ldp inter-area-fecs count
```

```
-----
Num. IPv4 FEC(s) : 9
-----
```

```
-----
Num. IPv6 FEC(s) : 0
-----
```

```
-----
Total Num. FEC(s): 9
-----
```

```
#show ldp inter-area-fecs ipv4 count
```

```
-----
Num. IPv4 FEC(s) : 9
-----
```

```
#show ldp inter-area-fecs ipv6 count
```

```
-----
Num. IPv6 FEC(s) : 0
-----
```

show ldp inter-area-fecs prefix

Use this command to show all LDP inter-area FECs by prefix.

Use parameter count to show FEC count for each prefix.

Command Syntax

```
show ldp inter-area-fecs prefix (A.B.C.D/M|X:X::X:X/M) count
```

Parameter

A.B.C.D/M

IP prefix <network>/<length>, e.g., 35.0.0.0/8

X:X::X:X/M

IPv6 prefix <network>/<length>, e.g., 3ffe::/16

count

Count of FECs

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 4.1.

Examples

```
#show ldp inter-area-fecs prefix 4.4.4.0/30
LSR codes      : E/N - LSR is egress/non-egress for this FEC,
                  L - LSR received a label for this FEC,
                  > - LSR will use this route for the FEC

FEC            Code   Session      Out Label      ELC      Nexthop Addr
Matching RIB prefix - 4.4.4.0/30
4.4.4.1/32      NL>    1.1.1.1      24970          No       12.1.1.1
                  NL>    3.3.3.3      24329          No       23.1.1.2
4.4.4.2/32      NL>    1.1.1.1      24971          No       12.1.1.1
                  NL>    3.3.3.3      24330          No       23.1.1.2
4.4.4.3/32      NL>    1.1.1.1      24972          No       12.1.1.1
                  NL>    3.3.3.3      24331          No       23.1.1.2

#show ldp inter-area-fecs prefix 4.4.4.0/30 count
Matching RIB prefix - 4.4.4.0/30
-----
Num. IPv4 FEC(s) : 3
-----

#show ldp inter-area-fecs prefix 3ffe::/16
LSR codes      : E/N - LSR is egress/non-egress for this FEC,
                  L - LSR received a label for this FEC,
                  > - LSR will use this route for the FEC

FEC            Code   Session      Out Label      ELC      Nexthop Addr

#show ldp inter-area-fecs prefix 3ffe::/16 count
```

show ldp interface

Table 15. show ldp fec output fields details

| Field | Description |
|--------------|--|
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| Session | Reports the current session state. |
| Out Label | Label received from downstream neighbor for route. |
| Nexthop addr | Displays the IP address of the next hop. |

Use this command to display the list of all interfaces on the current LSR, and to indicate whether a given interface is label-switching or not.

Command Syntax

```
show ldp interface
show ldp interface IFNAME
```

Parameter

IFNAME

Displays the name of the interface.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following output displays a list of all interfaces on the LSR.

```
#show ldp interface
Interface    LDP Identifier    Label-switching    Merge Capability
eth0 10.10.0.11:0 Disabled N/A
lo 10.10.0.11:0 Disabled N/A
eth1 10.10.0.11:0 Enabled Merge capable
eth2 10.10.0.11:0 Enabled Merge capable
vmnet1 10.10.0.11:0 Disabled N/A
```

The following is a sample output from the `show ldp interface IFNAME` command displaying information about the specified interface `eth1`.

```
#show ldp interface eth1
Status : Enabled
Primary IP Address : 192.168.3.4
Interface Type : Ethernet
Label Merge Capability : Merge Capable
Hello Interval : 5
Targeted Hello Interval : 15
```



```

Hold Time           : 15
Targeted Hold Time  : 45
Keepalive Interval  : 10
Keepalive Timeout   : 30
Advertisement Mode   : Downstream On Demand
Label Retention Mode : Liberal
Administrative Groups : myGroup

```

[Table 16](#) explains the show command output fields.

Table 16. show ldp interface output fields details

| Field | Description |
|--------------------------|---|
| Interface | Name of the interface. |
| LDP Identifier | LDP identifier for this protocol. |
| Label-switching | Status of the label-switching on interface .. |
| Merge Capability | Used to override the default merge capability setting of all the interfaces. |
| Status | Status of the ldp interface. |
| Primary IP Address | Address of the primary Internet protocol in the interface. |
| Interface Type | Type of interface. |
| Label Merge Capability | Used to override the default merge capability setting of all the interfaces for the label. |
| Hello Interval | Sets the interval for sending unicast hello packets to peers. |
| Targeted Hello Interval | Sets the interval for sending unicast hello packets to targeted peers. |
| Hold time | Sets the time-out value to peers. |
| Targeted Hold time | Sets the time-out value that is the time that the router waits before rejecting an adjacency with targeted peers. |
| Keepalive Interval | Used to set the interval for sending keep-alive messages to the peer in order to maintain a session. |
| Keepalive Timeout | Time-out value for rejecting a session with a peer. |
| Label Advertisement Mode | Used to set the label advertisement mode for an interface for the current LSR to either downstream-on-demand (label is sent only when requested) or downstream-unsolicited (label is sent unrequested). |
| Label Retention Mode | Used for all labels exchanged via the given interface. |
| Administrative Groups | Administrative group to be used for links. |

show ldp lsp

Use this command to display LDP LSP and, optionally, advertise-label information.

Command Syntax

```
show ldp lsp
show ldp lsp prefix detail
show ldp lsp (prefix|detail)
```

Parameter

prefix

Displays advertise-label information in addition to LDP LSP information.

detail

Displays advertise-label information in addition to LDP LSP information.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3. and was updated in OcNOS version 5.1.

Example

The following is a sample output from the `show ldp lsp prefix detail` command displaying LDP LSP prefix information with advertise-label information.

When LDP LFA FRR in not enabled:

```
#show ldp lsp
DOWNSTREAM LSP :
  FEC                Nexthop Addr      State      Label      Req.ID      Attr
  10.0.1.0/31         connected    Established impl-null  0
  10.0.1.0/31         connected    Established none      0          None
  10.0.2.0/31         connected    Established 24322      0
  10.0.2.0/31         10.0.1.1    Established impl-null  0
  10.0.3.0/31         connected    Established 24323      0
  10.0.3.0/31         10.0.4.1    Established impl-null  0
  10.0.4.0/31         connected    Established impl-null  0
  10.0.4.0/31         connected    Established none      0          None
  11.1.1.1/32         connected    Established none      0          None
  22.1.1.1/32         10.0.1.1    Established impl-null  0
  33.1.1.1/32         10.0.4.1    Established 24323      0
  33.1.1.1/32         10.0.1.1    Established 24324      0
  44.1.1.1/32         10.0.4.1    Established impl-null  0
  192.168.254.0/24    connected    Established impl-null  0
  192.168.254.0/24    connected    Established impl-null  0
  192.168.254.0/24    connected    Established none      0          None

UPSTREAM LSP :
  FEC                State      Label      Req.ID      Attr
  10.0.1.0/31         Established impl-null  0          None
  10.0.1.0/31         Established impl-null  0          None
```

```

10.0.2.0/31      Established      24320      0      None
10.0.3.0/31      Established      24321      0      None
10.0.4.0/31      Established      impl-null   0      None
10.0.4.0/31      Established      impl-null   0      None
11.1.1.1/32      Established      impl-null   0      None
11.1.1.1/32      Established      impl-null   0      None
22.1.1.1/32      Established      24322      0      None
44.1.1.1/32      Established      24324      0      None
192.168.254.0/24 Established      impl-null   0      None
192.168.254.0/24 Established      impl-null   0      None

```

When LDP LFA FRR in enabled:

```

#show ldp lsp
DOWNSTREAM LSP :
  FEC          Nexthop Addr      State      Label      Req.ID      Attr      Code
  10.0.1.0/31   connected    Established impl-null   0           None
  10.0.1.0/31   connected    Established none        0           None
  10.0.2.0/31   connected    Established 24322      0           0
  10.0.2.0/31   10.0.4.1     Established 24322      0           B
  10.0.2.0/31   10.0.1.1     Established impl-null   0           P
  10.0.3.0/31   connected    Established 24323      0
  10.0.3.0/31   10.0.1.1     Established 24323      0           B
  10.0.3.0/31   10.0.4.1     Established impl-null   0           P
  10.0.4.0/31   connected    Established impl-null   0
  10.0.4.0/31   connected    Established none        0           None
  11.1.1.1/32   connected    Established none        0           None
  22.1.1.1/32   10.0.1.1     Established impl-null   0           P
  33.1.1.1/32   10.0.4.1     Established 24323      0           P
  33.1.1.1/32   10.0.4.1     Established 24323      0           B
  33.1.1.1/32   10.0.1.1     Established 24324      0           P
  33.1.1.1/32   10.0.1.1     Established 24324      0           B
  44.1.1.1/32   10.0.4.1     Established impl-null   0
  192.168.254.0/24 connected    Established impl-null   0
  192.168.254.0/24 connected    Established impl-null   0
  192.168.254.0/24 connected    Established none        0           None

UPSTREAM LSP :
  FEC          State      Label      Req.ID      Attr
  10.0.1.0/31   Established impl-null   0           None
  10.0.1.0/31   Established impl-null   0           None
  10.0.2.0/31   Established 24320      0           None
  10.0.3.0/31   Established 24321      0           None
  10.0.4.0/31   Established impl-null   0           None
  10.0.4.0/31   Established impl-null   0           None
  11.1.1.1/32   Established impl-null   0           None
  11.1.1.1/32   Established impl-null   0           None
  22.1.1.1/32   Established 24322      0           None
  44.1.1.1/32   Established 24324      0           None
  192.168.254.0/24 Established impl-null   0           None
  192.168.254.0/24 Established impl-null   0           None

```

[Table 17](#) explains the show command output fields.

Table 17. show ldp lsp output fields details

| Field | Description |
|--------------|--|
| Session peer | Used to group and apply the configuration of general session commands to groups of neighbors that share common session configuration elements. |
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| Nexthop addr | Displays the IP address of the next hop. |

Table 17. show ldp lsp output fields details (continued)

| | |
|--------|---|
| State | Displays the current status of the ldp. |
| Label | Details of the ldp downstream labels. |
| Req.ID | Request identifier for the protocol. |
| Attr | The attribute is used to sent to a customer router. |
| Code | Show if an entry is principal or backup. |

show ldp mpls-l2-circuit

Use this command to display summarized Layer-2 Virtual Circuit information about all MPLS virtual circuits configured on the current LSR. When the Virtual Circuit ID is specified, this command displays summarized information for the Virtual Circuit matching the specified ID only.

Command Syntax

```
show ldp mpls-l2-circuit
show ldp mpls-l2-circuit <1-4294967295>
show ldp mpls-l2-circuit detail
show ldp mpls-l2-circuit count
show ldp mpls-l2-circuit <1-4294967295> detail
```

Parameter

<1-4294967295>

Indicates the virtual circuit ID.

detail

Displays detailed LDP information.

count

Count of PWs from LDP standpoint.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output of this command displaying summarized information of VID 1000:

```
#show ldp mpls-l2-circuit 1000
Transport Client VC Trans Local Remote Destination
VC ID Binding State Type VC Label VC Label Address
1000 eth2 UP ethernet 640 640 192.168.0.80

#show ldp mpls-l2-circuit
Transport Client VC Trans Local Remote Destination
VC ID Binding State Type VC Label VC Label Address
1000 eth2 UP ethernet 640 640 192.168.0.80
2000 eth3 UP ethernet 641 648 192.168.0.80
3000 eth4 UP ethernet 642 645 192.168.0.90
```

The following is a sample output of this command when using the `detail` parameter:

```
#show ldp mpls-l2-circuit detail
vcid: 100, type: ethernet, local groupid: 4, remote groupid: 4 (vc is up)
destination: 10.0.0.2, Peer LDP Ident: 10.0.0.2
Local label: 53120, remote label: 53120
Access IF: eth3, Network IF: eth4
Local MTU: 1500, Remote MTU: 1500
Local Control Word: 0, Remote Control Word: 0
```

```

Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status :
Not Forwarding
Remote PW Status :
Not Forwarding
Standby

```

[Table 18](#) explains the show command output fields.

Table 18. show ldp mpls-l2-circuit output fields details

| Field | Description |
|-----------------------------|---|
| Transport VC ID | Transport VC identifier for the protocol. |
| Client Binding | Show whether the interface is client bound and (if bound) with which client. |
| VC State | State of the VC. |
| Trans Type | Type of transmit. |
| Local VC Label | Incoming VC label details. |
| Remote VC Label | Outgoing VC label details. |
| Destination Address | Destination IP address for the protocol. |
| VCid | Address for the VC. |
| Type | Type of Ethernet interface. |
| local groupid | Address for the local group. |
| remote groupid | Address for the remote group. |
| destination | Destination IP address. |
| Peer LDP Ident | Identification for the peer LDP. |
| Local label | Number of Local label |
| remote label | Number remote label. |
| Access IF | Map the access port. |
| Network IF | Map the network port in the interface. |
| Local MTU | Number of local MTU., Remote MTU - Number of local MTU. |
| Local Control Word | Number of local control word. |
| Remote Control Word | Number of local control word. |
| Local PW Status Capability | PW Status capability of Local end of PW. |
| Remote PW Status Capability | PW Status capability of Remote end of PW. |
| Current PW Status TLV | A data structure used to encode optional information in a data communications protocol. |

Table 18. show ldp mpls-l2-circuit output fields details (continued)

| Field | Description |
|------------------|--------------------------------|
| Local PW Status | PW Status of Local end of PW. |
| Remote PW Status | PW Status of Remote end of PW. |

show ldp rlfa-routes

Use this command to display LDP remote LFA routes.

Command Syntax

```
show ldp rlfa-routes
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced in OcNOS version 5.1.

Example

```
#show ldp rlfa-routes
Fec          Primary-NH      Backup-NH      rLFA-Addr      Out-Intf  Outer-label  Inner-label
2.2.2.2       20.1.1.1      30.1.1.1      1.1.1.1        xe5       24320       24324
3.3.3.3       30.1.1.1      20.1.1.1      1.1.1.1        xe12      24320       24325
```

show ldp rlfa-routes count

Use this command to display ldp rlfa routes count.

Command Syntax

```
show ldp rlfa-routes count
```

Parameter

count

Counts the no of ldp-rlfa routes

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 6.6.1

Example

```
PE2#show ldp rlfa-routes count
-----
Total no of ldp-rlfa routes : 4
-----
```

show ldp routes

Use this command to display LDP routes.

Command Syntax

```
show ldp routes
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3 and was updated in OcNOS version 5.1.

Example

When LDP LFA FRR in not enabled:

```
#show ldp routes
Prefix Addr      Nexthop Addr    Intf
10.0.1.0/31      0.0.0.0         eth2
10.0.2.0/31      10.0.1.1        eth2
                  10.0.4.1        eth1
10.0.3.0/31      10.0.4.1        eth1
                  10.0.1.1        eth2
10.0.4.0/31      0.0.0.0         eth1
11.1.1.1/32      0.0.0.0         lo
22.1.1.1/32      10.0.1.1        eth2
33.1.1.1/32      10.0.1.1        eth2
                  10.0.4.1        eth1
44.1.1.1/32      10.0.4.1        eth1
192.168.254.0/24 0.0.0.0         eth0
```

When LDP LFA FRR in enabled:

```
#show ldp routes
Prefix Addr      Nexthop Addr    Intf      Backup Addr    Backup Intf
10.0.1.0/31      0.0.0.0         eth2      -              -
10.0.2.0/31      10.0.1.1        eth2      10.0.4.1       eth1
10.0.3.0/31      10.0.4.1        eth1      10.0.1.1       eth2
10.0.4.0/31      0.0.0.0         eth1      -              -
11.1.1.1/32      0.0.0.0         lo        -              -
22.1.1.1/32      10.0.1.1        eth2      -              -
33.1.1.1/32      10.0.1.1        eth2      10.0.4.1       eth1
                  10.0.4.1        eth1      10.0.1.1       eth2
44.1.1.1/32      10.0.4.1        eth1      -              -
192.168.254.0/24 0.0.0.0         eth0      -              -
```

[Table 19](#) explains the show command output fields.

Table 19. show ldp routes output fields details

| | |
|--------------|---|
| Prefix Addr | Details of the network address prefix. |
| Nexthop Addr | Displays the IP address of the next hop. |
| Intf | Displays an interface name. |
| Backup Addr | Displays the IP address of the backup next hop. |
| Backup Intf | Displays a backup interface name. |

show ldp routes count

Use this command to display the ldp route count.

Command Syntax

```
show ldp routes count
```

Parameter

count

Counts the no of ldp routes

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 6.6.1

Example

```
PE2#show ldp routes count
```

```
-----  
Total no of LDP routes : 8  
-----
```

show ldp session

Use this command to display sessions established between this LSR and other LSRs.

Command Syntax

```
show ldp session
show ldp session A.B.C.D
show ldp session X:X::X:X
show ldp session detail
show mpls ldp session
show mpls ldp session A.B.C.D
show mpls ldp session X:X::X:X
show mpls ldp session detail
```

Parameter

A.B.C.D

IPv4 address of the peer.

X:X::X:X

IPv6 address of the peer.

detail

Displays the detailed output for all the existing LDP sessions.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The below example displays the output for operational LDP session:

```
OcNOS#show ldp session 192.168.3.5
Session state: OPERATIONAL
Session role: Passive
TCP Connection: Established
IP Address for TCP: 192.168.3.5
Interface being used : eth1
Peer LDP ID: 10.10.0.18:0
Last Down Reason: None
Peer Password : mypwd
Authentication type: MD5
Adjacencies: 192.168.3.5
192.168.4.5
Advertisement mode: Downstream Unsolicited
Label retention mode : Liberal
Graceful Restart : Capable
Reconnect Timeout : 120
Recovery Timeout (max) : 120
Recovery Timeout [negotiated] : 0 [120]
Keepalive Timeout: 30
Reconnect Interval: 15
Address List received : 192.168.3.5
```

```

192.168.4.5
Received Labels :FecLabelMaps To
IPv4:10.10.0.0/24
impl-null none
IPv4:192.168.3.0/24 impl-null none
IPv4:192.168.4.0/24 impl-null none
IPv4:192.168.5.0/24 impl-null none
Sent Labels :FecLabelMaps To
IPv4:10.10.0.0/24
impl-null none
IPv4:192.168.3.0/24 impl-null none
IPv4:192.168.4.0/24 impl-null none

```

The below example displays the output for non existent LDP session:

```

OcNOS#show ldp sess detail
Peer IP Address: 2.2.2.2

Session state           : NON_EXISTENT
Session role           : Passive
TCP Connection          : Not Established
IP Address for TCP      : 2.2.2.2
Interface being used    : xe10
Peer LDP ID             : 2.2.2.2:0
Last Down Reason        : Command Line Interface (CLI) Command
Flap Time               : 00:00:07
Session Down Reason     : KeepAlive Timer Expired
Preferred Peer LDP Password : Not Set
Adjacencies             : 10.10.10.2
Advertisement mode       : Downstream Unsolicited
Label retention mode     : Liberal
Graceful Restart        : Not Capable
Keepalive Timeout       : 3
Reconnect Interval      : 15
Received Labels :      Fec      Label      Maps To
Sent Labels :      Fec      Label      Maps To

```

[Table 20](#) explains the show command output fields.

Table 20. show ldp session output fields details

| Field | Description |
|---------------------|---|
| Session state | Reports the current session state. |
| Session role | Displays the status of the session role. |
| TCP Connection | Details of the TCP connection. |
| IP Address for TCP | Transmission control protocol IP address for the network. |
| Interface | Name of interface used in the network. |
| Peer LDP ID | Identifier for the peer LDP. |
| Last Down Reason | Displays the reason for the initial failure of the LDP session. |
| Flap Time | Displays the duration of the time since the failure of the LDP session. |
| Session Down Reason | Displays the reason for the ongoing failure of the LDP session. |
| Peer Password | Credential details for the neighbor. |
| Authentication type | Type of authentication. |
| Adjacencies | IP address for the neighbor adjacencies. |

Table 20. show ldp session output fields details (continued)

| Field | Description |
|-------------------------------|---|
| Advertisement mode | Details of the advertisement mode. |
| Label retention mode | Details of the label retention mode. |
| Graceful Restart | Indicates if the peer session is “Capable” or “Not Capable”. |
| Reconnect Timeout | The amount of time the router keeps the labels until session re-connection, the value is the lower value between local and remote neighbor-liveness timer. It appears when the session is GR capable. |
| Recovery Timeout (max) | Indicates the amount of time for the recovery session to send the initialization message to the peer, according to the local max-recovery timer. It appears when the session is GR capable. |
| Recovery Timeout [negotiated] | Indicates the actual timer value and the initial amount of time to recovery session (between brackets) that is negotiated with the peer to the lower value between local and remote values. Negotiated value 0 indicates the labels are not preserved after session disconnection. It appears when the session is GR capable. |
| Keepalive Interval | Used to set the interval for sending keep-alive messages to the peer in order to maintain a session. |
| Keepalive Timeout | Time-out value for rejecting a session with a peer. |
| Address List received | List of address that is received from neighbor. |
| Received Labels | Number of labels received from neighbor session. |
| Sent Labels | Number of labels transmitted to neighbor session. |

show ldp session all count

Use this command to display the total number of UP and Down LDP sessions.

Command Syntax

```
show ldp session all count
```

Parameter

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

This command displays the number of up and down multicast LDP sessions

```
OcNOS#show ldp session all count
```

```
-----  
  
session up count    : 2  
session down count  : 0  
total count         : 2  
  
-----
```

[Table 21](#) explains the show command output fields.

Table 21. show ldp session output fields details

| Field | Description |
|--------------------|--|
| session up count | Displays the number of running LDP sessions. |
| session down count | Displays the number of failed LDP sessions. |
| total count | Displays the total number of LDP sessions. |

show ldp session multicast count

Use this command to display the number of up and down multicast LDP sessions.

Command Syntax

```
show ldp session multicast count
```

Parameter

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

This command displays the number of up and down multicast LDP sessions

```
OcNOS#show ldp session multicast count
-----
session up count      : 2
session down count    : 0
total count           : 2
```

[Table 22](#) explains the show command output fields.

Table 22. show ldp session output fields details

| Field | Description |
|--------------------|--|
| session up count | Displays the number of running multicast LDP sessions. |
| session down count | Displays the number of failed multicast LDP sessions. |
| total count | Displays the total number of multicast LDP sessions. |

show ldp session targeted count

Use this command to display the number of UP and Down targeted LDP sessions.

Command Syntax

```
show ldp session targeted count
```

Parameter

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

This command displays the number of up and down multicast LDP sessions

```
OcNOS#show ldp session targeted count

-----

session up count      : 2
session down count    : 0
total count           : 2

-----
```

[Table 23](#) explains the show command output fields.

Table 23. show ldp session output fields details

| Field | Description |
|--------------------|---|
| session up count | Displays the number of running targeted LDP sessions. |
| session down count | Displays the number of failed targeted LDP sessions. |
| total count | Displays the total number of targeted LDP sessions. |

show ldp statistics

Use this command to display LDP statistics.

Command Syntax

```
show ldp statistics
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output from the show ldp statistics command.

```
#show ldp statistics

=====
LSR ID = 180.1.1.1:0 : INTERFACE NAME: xe8
=====
PacketType                Total
                          Sent      Received
Notification                0          0
Hello                       0          0
Initialization              0          0
Keepalive                   0          0
Address                     0          0
Address Withdraw            0          0
Label Mapping               0          0
Label Request               0          0
Label Withdraw              0          0
Label Release               0          0
Request Abort               0          0

Session EventType          Count
-----
Session Attempts failed    0
No Hello Errors             0
Advertisement Mode Errors   0
Max Pdu Length Errors      0
Label Range Errors          0
Bad LDP Identifier Errors   0
Bad PDU Length Errors       0
Bad Message Length Errors   0
Bad TLV Length Errors       0
Malformed TLV Value Errors  0
Keep Alive Timer Expired Errors 0
+ Unknown Fec Errors        0
Shutdown Notifications received 0
Shutdown Notifications sent  0
=====
```

[Table 24](#) explains the show command output fields.

Table 24. show ldp statistics output fields details

| Field | Description |
|---------------|--|
| LSR ID | Identifier of the LSR. |
| Targeted Peer | Targeted LDP neighbor can improve the label convergence time compared to the convergence time with directly connected LDP peers when there are flapping links. |
| Packet Type | Type of packet in the interface that has been received or transmitted to the neighbors. |
| Total | Number of total packets that has been received and transmitted. |

show ldp statistics advertise-labels

Use this command to display the count per each operation filtered by an advertisement list.

Command Syntax

```
show ldp statistics advertise-labels
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output from the `show ldp statistics advertise-labels` command.

```
#show ldp statistics advertise-labels
Advertisement spec:
  Prefix list = pfx1; Peer plist = pfx1
  Deny : Label Mapping = 2
         Label Request = 0
Prevent the distribution of any assigned labels
  Deny : Label Mapping = 9
         Label Request = 3
#
```

[Table 25](#) explains the show command output fields.

Table 25. show ldp statistics advertise-labels output fields details

| Field | Description |
|--------------------|--|
| Advertisement spec | Details of the advertisement spec. |
| Prefix list | It is an ordered list and entries are evaluated in order of increasing sequence number. |
| Peer plist | The peer keyword enables the device to receive time requests and used to synchronize itself to the servers specified in the access list. |
| Label Mapping | Number of label mapping that is denied. |
| Label Request | Number of label request that is denied. |

show ldp targeted-peers

Use this command to display the list of targeted peers configured on the current LSR.

Command Syntax

```
show ldp targeted-peers
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output from the `show ldp targeted-peers` command.

```
#show ldp targeted-peers
IP Address      Interface
192.168.201.2   eth1
```

[Table 26](#) explains the show command output fields.

Table 26. show ldp targeted-peers output fields details

| Field | Description |
|------------|--|
| IP Address | Internet protocol address for the interface. |
| Interface | Name of the interface. |

show ldp upstream

Use this command to display the status of all upstream sessions and label information exchanged.

Command Syntax

```
show ldp upstream (session A.B.C.D|prefix A.B.C.D/M|)
```

Parameter

prefix

Ldp upstream for prefix

session

LDP Session

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3. and was updated in OcNOS version 6.6.1.

Example

The following is a sample output of the `show ldp upstream` command showing the status of all upstream sessions.

```
#show ldp upstream
Session peer 1.1.1.1:
  FEC          State      Label    Req.ID  Attr
  60.0.0.0/24   Established 52486    0       None
  4.4.4.4/32    Established 52484    0       None
  50.0.0.0/24   Established 52483    0       None
  40.0.0.0/24   Established impl-null 0       None
  30.0.0.0/24   Established impl-null 0       None
  20.0.0.0/24   Established impl-null 0       None
  10.0.2.0/24   Established impl-null 0       None
  5.5.5.5/32    Established 52482    0       None
  3.3.3.3/32    Established 52481    0       None
  2.2.2.2/32    Established impl-null 0       None
Session peer 3.3.3.3:
  FEC          State      Label    Req.ID  Attr
  60.0.0.0/24   Established 52487    0       None
  4.4.4.4/32    Established 52485    0       None
  1.1.1.1/32    Established 52480    0       None
  40.0.0.0/24   Established impl-null 0       None
  30.0.0.0/24   Established impl-null 0       None
  20.0.0.0/24   Established impl-null 0       None
  10.0.2.0/24   Established impl-null 0       None
  2.2.2.2/32    Established impl-null 0       None
Session peer 4.4.4.4:
  FEC          State      Label    Req.ID  Attr
  50.0.0.0/24   Established 52483    0       None
  40.0.0.0/24   Established impl-null 0       None
  30.0.0.0/24   Established impl-null 0       None
  20.0.0.0/24   Established impl-null 0       None
  10.0.2.0/24   Established impl-null 0       None
```

| | | | | |
|------------|-------------|-----------|---|------|
| 3.3.3.3/32 | Established | 52481 | 0 | None |
| 2.2.2.2/32 | Established | impl-null | 0 | None |
| 1.1.1.1/32 | Established | 52480 | 0 | None |

```

R1#show ldp upstream session 11.11.11.11
! [execution timestamp : 2025 Apr 03 10:58:55]
FEC State Label Req.ID Attr Installed Refcnt
Session peer 11.11.11.11:
14.14.14.14/32 Established 26249 0 None Yes 6
10.1.2.0/30 Established 26288 0 None Yes 4
[Summary] total upstreams: 2
[state] depend: 0, idle: 0, resrc-pre: 0, resp: 0, resrc: 0, estab: 2, rel: 0
[label] allocated: 2, impl-null: 0, expl-null: 0, invalid: 0, installed: 0, common-install: 2

R1#show ldp upstream prefix 14.14.14.14/32
! [execution timestamp : 2025 Apr 03 10:59:03]
FEC State Label Req.ID Attr Installed Refcnt
Session peer 11.11.11.11:
14.14.14.14/32 Established 26249 0 None Yes 6
Session peer 14.14.14.14:
14.14.14.14/32 Established 26249 0 None skipped 6

```

The following table explains the show command output fields.

Table 27. show ldp upstream output fields details

| Field | Description |
|--------------|--|
| Session peer | Details of the session peers. |
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| State | Reports the current session state. |
| Label | Number of Label received from upstream neighbor for route. |
| Req.ID | Requested session identifier for the protocol. |
| Attr | The attribute is used to sent packets to a customer router. |

show ldp downstream

Use this command to display the status of all downstream sessions and the label information exchanged.

Command Syntax

```
show ldp downstream (session A.B.C.D|prefix A.B.C.D/M|)
```

Parameter

prefix

ldp upstream for prefix

session

LDP Session

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3 and was updated in OcNOS version 6.6.1.

Example

The following is an output from the `show ldp downstream` command showing the status of all downstream sessions.

When LDP LFA FRR in not enabled:

```
#show ldp downstream
Session peer 44.1.1.1:
  FEC           Nexthop Addr   State           Label           Req.ID          Attr
  192.168.254.0/24 connected      Established      impl-null        0
  33.1.1.1/32    10.0.4.1      Established      24323            0
  10.0.4.0/31    connected      Established      impl-null        0
  10.0.2.0/31    connected      Established      24322            0
  44.1.1.1/32    10.0.4.1      Established      impl-null        0
  10.0.3.0/31    10.0.4.1      Established      impl-null        0
Session peer 22.1.1.1:
  FEC           Nexthop Addr   State           Label           Req.ID          Attr
  192.168.254.0/24 connected      Established      impl-null        0
  10.0.3.0/31    connected      Established      24323            0
  10.0.1.0/31    connected      Established      impl-null        0
  33.1.1.1/32    10.0.1.1      Established      24324            0
  22.1.1.1/32    10.0.1.1      Established      impl-null        0
  10.0.2.0/31    10.0.1.1      Established      impl-null        0
```

When LDP LFA FRR in enabled:

```
#show ldp downstream
Codes: P - Primary route, B - Backup route
Session peer 44.1.1.1:
  FEC           Nexthop Addr   State           Label           Req.ID          Attr          Code
  192.168.254.0/24 connected      Established      impl-null        0
  33.1.1.1/32    10.0.4.1      Established      24323            0                                     P
```

```

33.1.1.1/32      10.0.4.1      Established    24323      0      B
10.0.4.0/31      connected     Established    impl-null   0
10.0.2.0/31      connected     Established    24322      0
10.0.2.0/31      10.0.4.1      Established    24322      0      B
44.1.1.1/32      10.0.4.1      Established    impl-null   0      P
10.0.3.0/31      10.0.4.1      Established    impl-null   0      P

```

Codes: P - Primary route, B - Backup route

Session peer 22.1.1.1:

| FEC | Nexthop Addr | State | Label | Req.ID | Attr | Code |
|------------------|--------------|-------------|-----------|--------|------|------|
| 192.168.254.0/24 | connected | Established | impl-null | 0 | | |
| 10.0.3.0/31 | connected | Established | 24323 | 0 | | |
| 10.0.3.0/31 | 10.0.1.1 | Established | 24323 | 0 | | B |
| 10.0.1.0/31 | connected | Established | impl-null | 0 | | |
| 33.1.1.1/32 | 10.0.1.1 | Established | 24324 | 0 | | P |
| 33.1.1.1/32 | 10.0.1.1 | Established | 24324 | 0 | | B |
| 22.1.1.1/32 | 10.0.1.1 | Established | impl-null | 0 | | P |
| 10.0.2.0/31 | 10.0.1.1 | Established | impl-null | 0 | | P |

```

R1#show ldp downstream session 11.11.11.11
! [execution timestamp : 2025 Apr 03 10:59:25]
FEC Nexthop Addr State Label Req.ID Attr Installed Code

Codes: P - Primary route, B - Backdown route
Session peer 11.11.11.11:
11.11.11.11/32 non-nh Established 26276 0 None No -
14.14.14.14/32 non-nh Established 26241 0 None No -
[Summary] total downstreams: 3
[state] depend: 0, idle: 0, resp: 0, estab: 2, inactive: 0
[label] user-label: 2, impl-null: 0, expl-null: 0, invalid: 0

```

```

R1#show ldp downstream prefix 14.14.14.14/32
! [execution timestamp : 2025 Apr 03 10:59:34]
FEC Nexthop Addr State Label Req.ID Attr Installed Code

Codes: P - Primary route, B - Backdown route
Session peer 11.11.11.11 :
14.14.14.14/32 non-nh Established 26241 0 None No -
Session peer 14.14.14.14 :
14.14.14.14/32 60.60.60.1 Established impl-null 0 None Yes

```

[Table 28](#) explains the show command output fields.

Table 28. show ldp downstream output fields details

| Field | Description |
|--------------|--|
| Session peer | Used to group and apply the configuration of general session commands to groups of neighbors that share common session configuration elements. |
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| Nexthop addr | Displays the IP address of the next hop. |
| State | Displays the current status of the ldp. |
| Label | Details of the ldp downstream labels. |
| Req.ID | Request identifier for the protocol. |
| Attr | The attribute is used to sent to a customer router. |
| Code | Show if an entry is principal or backup. |

show ldp vpls

Use this command to display information about all VPLS instances. Specify the VPLS ID to display information about a specific VPLS instance.

Command Syntax

```
show ldp vpls <1-4294967295> (count|)
show ldp vpls count
show ldp vpls detail
show ldp vpls (no-vc|)
```

Parameter

<1-4294967295>

Display the VPLS identifier.

count

Display VPLS count from LDP standpoint.

detail

Display detailed LDP VPLS information.

no-vc

Specify not display L2VC information.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output of the `show ldp vpls` command displaying information about all VPLS instances.

```
#show ldp vpls
VPLS-ID    Peer Address      State  Type      Label-Sent  Label-Rcvd
1          192.168.0.80      Up     vpls      16          640
1          192.168.0.90      Up     vpls      18          642
2          192.168.0.80      Up     vpls      19          641
2          192.168.0.90      Up     vpls      17          643
```

The following is an output of the `show ldp vpls detail` command PW is Up:

```
VPLS Identifier      : 1
Peer IP              : 110.1.1.1
VC State             : UP
VC Type              : ethernet
VC Label Sent        : 24960
VC Label Received    : 24960
Local MTU            : 9000
Remote MTU           : 9000
+ Local FEC Type      : 128
+ Remote FEC Type     : 128
```

```

Local Control Word : enabled
Remote Control Word: disabled
Current use       : disabled
Local Flow Label  : disabled
Remote Flow Label : disabled
Current use       : disabled
Local PW Status Capability : disabled
Remote PW Status Capability : enabled
Current PW Status TLV : disabled
LDP-VPLS Signaled Time : 00:03:05

```

The following is an output of the `show ldp vpls detail` command PW is Down:

```

VPLS Identifier      : 1
Peer IP              : 110.1.1.1
VC State             : DOWN
VC Type              : ethernet
VC Label Sent        : 24961
VC Label Received    : 24961
Local MTU            : 1500
Remote MTU           : 9000
(Error: MTU mismatched)
+ Local FEC Type     : 128
+ Remote FEC Type    : 129
+ (Error: FEC Type mismatched)
Local Control Word : enabled
Remote Control Word: disabled
Current use       : disabled
Local Flow Label  : disabled
Remote Flow Label : disabled
Current use       : disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
LDP-VPLS Signaled Time : -

* Label released due to attribute mismatch
- MTU mismatch
+ - FEC Type mismatch

```

The following is a sample output of `show ldp vpls count` displaying information about total, active and inactive VPLS instances from LDP.

```

#show ldp vpls count
-----
Total VPLS instances      : 2
Active VPLS instances     : 2
Inactive VPLS instances   : 0
-----

```

[Table 29](#) explains the show command output fields.

Table 29. show ldp vpls output fields details

| Field | Description |
|------------|--|
| VPLS-ID | Identification details of the VPLS. |
| Peer Addr | IP address of the peer device. |
| State | Reports the current session state. |
| Type | Type of protocol in network. |
| Label-Sent | Number of packets transmitted to neighbor. |

Table 29. show ldp vpls output fields details (continued)

| Field | Description |
|-------------------------|--|
| Label-Rcvd | Number of packets received from neighbor. |
| Total VPLS instances | Number of total VPLS instance in the protocol. |
| Active VPLS instances | Number of active VPLS instance. |
| Inactive VPLS instances | Number of inactive VPLS instance. |

show mpls ldp discovery

Use this command to display the sources for locally generated LDP Discovery Hello PDUs, and to indicate whether an interface is label-switching.

Command Syntax

```
show mpls ldp discovery
show mpls ldp discovery IFNAME
```

Parameter

IFNAME

Interface name.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls ldp discovery
Interface      LDP Identifier      Label-switching      Merge Capability
eth0           10.10.0.11:0        Disabled             N/A
lo             10.10.0.11:0        Disabled             N/A
eth1           10.10.0.11:0        Enabled              Merge capable
eth2           10.10.0.11:0        Enabled              Merge capable
vmnet1         10.10.0.11:0        Disabled             N/A
```

[Table 30](#) explains the show command output fields.

Table 30. show ldp discovery output fields details

| Field | Description |
|------------------|--|
| Interface | Name of the interface. |
| LDP Identifier | LDP identifier for this protocol. |
| Label-switching | Status of the label-switching on interface. |
| Merge Capability | Used to override the default merge capability setting of all the interfaces. |

show mpls ldp neighbor

Use this command to display LDP neighbor information.

Command Syntax

```
show mpls ldp neighbor
show mpls ldp neighbor detail
```

Parameter

detail

Details for adjacencies.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls ldp neighbor detail
IP Address  Interface Name  Holdtime  LDP ID
192.168.3.5  eth1             15       10.10.0.18:0
192.168.4.5  eth2             15       10.10.0.18:0
```

[Table 31](#) explains the show command output fields.

Table 31. show mpls ldp neighbor output fields

| Field | Description |
|----------------|--|
| IP Address | Address of the interface. |
| Interface Name | Name of the interface. |
| Holdtime | The amount of time this device waits between SPF. |
| LDP ID | Local label space ID. The first four bytes of an LDP ID is a platform IP address called the LDP router ID. The last two bytes are called the local label space ID. |

show mpls ldp parameter

Use this command to display LDP configuration parameters.

Command Syntax

```
show mpls ldp parameter
```

Parameter

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls ldp parameter
Router ID           : 0.0.0.0
LDP Version         : 1
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode  : Liberal
Label Control Mode    : Independent
Instance Loop Detection : Off
Request Retry         : Off
Propagate Release     : Disabled
Graceful Restart      : Disabled
Hello Interval        : 5
Targeted Hello Interval : 15
Hold time             : 15
Targeted Hold time    : 45
Keepalive Interval    : 10
Keepalive Timeout     : 30
Request retry Timeout  : 5
Transport Address data :
  Labelspace 0        : 192.168.201.2 (not in use)
Import BGP routes     : No
```

[Table 32](#) explains the show command output fields.

Table 32. show mpls ldp parameters output fields

| Field | Description |
|--------------------------|--|
| Router ID | A preferred interface address for LDP router. |
| LDP Version | Latest LDP version details. |
| Global Merge Capability | Override the default merge capability setting of all the interfaces. |
| Label Advertisement mode | Label advertisement mode details in the interface. |

Table 32. show mpls ldp parameters output fields (continued)

| Field | Description |
|-------------------------|---|
| Label retention mode | Label retention mode details in the interface. |
| Label Control Mode | Controls the mode used for handling label binding requests on interfaces. |
| Instance Loop Detection | Disables the LDP optional loop detection mechanism. |
| Request Retry | Request causes the target peer to respond with targeted Hello messages. |
| Propagate Release | Propagate release is disabled in the interface. |
| Graceful Restart | Graceful Restart (GR) is a mechanisms to prevent routing protocol re-convergence during a processor switchover. Hello Interval - Hello interval sets the interval for sending unicast hello packets to peers. |
| Targeted Hello Interval | Targeted hello interval sets the interval for sending unicast hello packets to targeted peers. |
| Hold time | Hold time sets the time-out value to peers. |
| Targeted Hold time | Time-out value is the time that the router waits before rejecting an adjacency with targeted peers. |
| Keepalive Interval | Keepalive interval sets the interval for sending keep-alive messages to the peer in order to maintain a session. |
| Keepalive Timeout | Time-out value for rejecting a session with a peer. |
| Request retry Timeout | Request for the maximum retry duration (the number of retries times the length of the timeout). |
| Transport Address data | Transport address advertised in LDP Discovery Hello messages sent on an interface. |
| Label space | Label used in a label binding is allocated from a set of possible labels called a label space. |
| Import BGP routes | The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding (VRF) instance table using an import route map. |

show ldp tunneling-fec

Use this command to display LDP FECs which are mapped to TE-tunnel for tunneling LDP over RSVP.

Command Syntax

```
show ldp tunneling-fec
```

Parameters

None

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 6.5.3.

Examples

This example shows **show ldp tunneling-fec** for upstream peer.

```
(OcNOS)# show ldp tunneling-fec
```

| FEC label | Tunnel-name | Tunnel-endpoint | Upstream-Peer | In-label | Out- |
|----------------|-------------|-----------------|---------------|----------|-------|
| ===== | | | | | |
| 10.12.183.5/32 | PE2-PE4 | 10.12.183.4/32 | 10.12.183.1 | 25029 | 24989 |
| 10.12.183.6/32 | PE2-PE4 | 10.12.183.4/32 | 10.12.183.1 | 25030 | 24990 |

show ldp tunneling

Use this command to display TE-tunnels enabled for tunneling LDP FECs.

Command Syntax

```
show ldp tunneling (TRUNKNAME|)
```

Parameters

TRUNKNAME

The trunk name is an optional parameter and represents the name of the trunk for which the details need to be displayed.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 6.5.3.

Examples

This example shows **show ldp tunneling** for upstream peer.

```
(OcNOS)#show ldp tunneling

Tunnel Name       : PE2-PE4
Tunnel Endpoint   : 10.12.183.4/32
Tunnel Cost       : 10
Tunnel Owner      : RSVP
Tunnel Status     : Up

FEC                Upstream-Peer   In-Label   Out-Label
=====
10.12.183.5/32     10.12.183.1   25029     24989
10.12.183.6/32     10.12.183.1   25030     24990

Total FEC tunneld by PE2-PE4 : 2
```

show ldp tunneling-tunnels

Use this command to display all TE-tunnels which are enabled for tunneling LDP FECs.

Command Syntax

```
show ldp tunneling-tunnels
```

Parameters

None

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 6.5.3.

Examples

This example shows **show ldp tunneling-tunnels** for upstream peer.

```
(OcNOS)# show ldp tunneling-tunnels
```

| Tunnel-name | Tunnel-endpoint | Status | Cost |
|-------------|-----------------|--------|------|
| ===== | | | |
| PE2-PE4 | 10.12.183.4/32 | Up | 10 |
| PE2-PE5 | 10.12.183.5/32 | Up | 20 |

RESOURCE RESERVATION PROTOCOL CONFIGURATION

| | |
|--|-----|
| RSVP-TE Configuration | 643 |
| RSVP-TE Overview | 643 |
| Configure RSVP-TE | 644 |
| Configure Global Parameters | 674 |
| MPLS RSVP PING and TRACEROUTE | 675 |
| MPLS RSVP Entropy Label Capabilities | 675 |
| RSVP-TE Facility Backup (Facility Bypass) | 677 |
| Topology | 677 |
| Configuration | 677 |
| Validation | 682 |
| Facility Bypass with Ring Topology Configuration | 688 |
| Configuration | 689 |
| Validation | 695 |
| RSVP-Multipath Configuration | 704 |
| Topology | 704 |
| Configuration | 704 |
| Validation | 710 |
| RSVP-TE Dynamic Facility Backup LSP (RSVP Auto Bypass) | 721 |
| Overview | 721 |
| Prerequisites | 721 |
| Configuration for RSVP Auto Bypass | 723 |
| RSVP Auto Bypass Commands | 737 |
| RSVP Graceful Restart Configuration | 752 |
| Overview | 752 |
| Configuration | 752 |
| Validation | 756 |
| Fast Reroute Configuration (One-to-one Method) | 758 |
| RTR-29 Configuration | 758 |
| RTR-43 Configuration | 761 |
| RTR-42 Configuration | 762 |
| RTR-41 Configuration | 764 |
| RTR-31 Configuration | 765 |
| RTR-30 Configuration | 766 |
| Validation | 768 |
| RSVP Detour Over Ring Topology | 771 |

| | |
|---|-----|
| Auto-Bandwidth with RSVP-TE | 790 |
| Overview | 790 |
| Prerequisites | 791 |
| Configuration for RSVP Auto-Bandwidth | 793 |
| Commands for RSVP Auto-Bandwidth | 801 |
| RSVP LSP Reoptimization | 834 |
| Overview | 834 |
| Characteristics | 834 |
| Prerequisites | 834 |
| Configuration | 835 |
| Running Configuration | 840 |
| Validation | 844 |
| Implementation Examples | 861 |
| Revised Commands | 861 |
| Troubleshooting | 861 |
| Glossary | 862 |

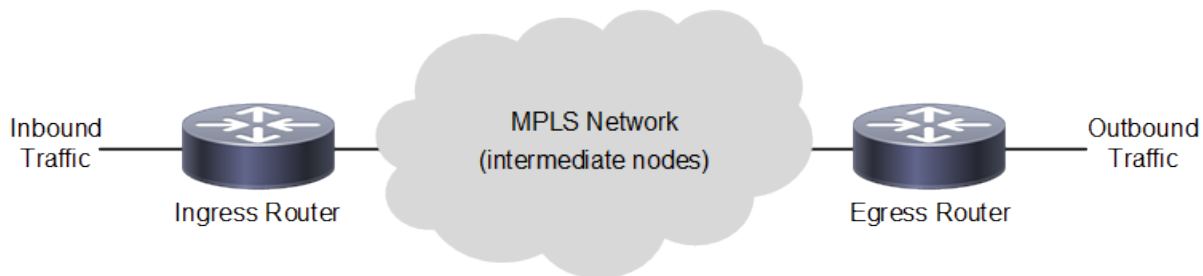
RSVP-TE Configuration

This chapter contains configurations for Resource Reservation Protocol - Traffic Engineering (RSVP-TE).

RSVP-TE Overview

RSVP-TE is a signaling protocol that supports explicit routing capabilities. To do this, an Explicit Route (ER) object is incorporated into RSVP PATH messages. This object encapsulates a sequence of hops that constitute the explicitly-routed path. Using the ER object, the paths taken by label-switched RSVP-MPLS flows can be pre-determined without conventional IP routing. An ER path can be administratively specified or computed based on CSPF and any policy requirements dictated by the operator through the trunk node, taking the current network state into consideration. A useful application of explicit routing is Traffic Engineering (TE). Using explicitly-routed LSPs, an ingress node can control the path through which traffic flows from itself, through the MPLS network, to the egress node. Explicit routing is therefore useful for the optimization of network resources and an increase in the quality of traffic-oriented performance.

Figure 21. Basic RSVP-TE Topology



RSVP-TE Architecture

RSVP-TE is a signaling protocol that supports explicit routing capabilities to establish LSPs in a MPLS network. OcNOS RSVP-TE:

- Creates explicitly-routed paths, which might not agree with the route suggested by the IGP (OSPF, RIP) being used. Explicitly-routed LSPs, by definition, do not follow the paths suggested by IGPs.
- Queries CSPF for a complete, end-to-end, explicit route based on constraints specified by the operator using RSVP commands.
- Performs make-before-break type re-routing of tunnels. (Make-before-break is the creation of a new LSP before the old one is torn down).
- Exchanges Hello messages to make node failures easier to detect. This means when there is no hello exchange between routers, then other node is assumed dead or offline (except in the case when the peer is known to not support Hellos).
- Provides statistical information of RSVP messages exchanged.

In addition, OcNOS RSVP-TE may be used in unison with BGP to generate MPLS/BGP VPNs, and in unison with LDP to generate Layer-2 Virtual Circuits.

Configure RSVP-TE



Note: The following configuration for establishing a trunk is required on all routers participating in label-switching. Based on the assumption that minimal configurations exist on all participating routers, other examples do not repeat this configuration.

Enable Label Switching - Minimal Configuration


To establish a trunk on a system:

1. Enable label-switching and RSVP-TE on all participating interfaces.
2. Configure a trunk on the ingress router to use the best available IGP path.

In this example, the Label Switched Path (LSP) is configured using minimal configuration and is setup using the best IP nexthop available. Each router along the path is chosen by the previous router by looking up the best nexthop available in its IP routing table.

Topology

Figure 22. Topology for Minimal Configuration



PE1 - NSM

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 192.168.0.63/32 secondary | Set the IP address for the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 10.10.23.63/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#commit | Commit the transaction. |

PE1 - RSVP-TE

| | |
|-------------------------|---|
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#commit | Commit the transaction. |

PE1 - OSPF

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the Routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 192.168.0.63 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 10.10.23.0/24 area 0 | Define the network (10.10.23.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 192.168.0.63/32 area 0 | Set the IP address of the loopback interface to 192.168.0.63/32. |
| (config-router)#commit | Commit the transaction. |

P - NSM

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 192.168.0.60/32 secondary | Set the IP address for the interface. |
| (config-if)#exit | Enable label switching on interface lo. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 10.10.23.60/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ip address 10.10.21.60/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#commit | Commit the transaction. |

P - RSVP-TE

| | |
|-------------------------|---|
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#commit | Commit the transaction. |

P - OSPF

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the Routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 192.168.0.60 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 10.10.23.0/24 area 0 | Define the first network (10.10.23.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 10.10.21.0/24 area 0 | Define the second network (10.10.21.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 192.168.0.60/32 area 0 | Set the IP address of the loopback interface to 192.168.0.63/32. |
| (config-router)#commit | Commit the transaction. |

PE2 - NSM

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 192.168.0.90/32 secondary | Set the IP address for the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ip address 10.10.21.90/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#commit | Commit the transaction. |

PE2 - RSVP-TE

| | |
|-------------------------|---|
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#commit | Commit the transaction. |

PE2 - OSPF

| | |
|--------------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the Routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing |

| | |
|---|--|
| | process. |
| <code>(config-router)#router-id 192.168.0.90</code> | Configure OSPF router-ID same as loopback interface IP address |
| <code>(config-router)#network 10.10.21.0/24 area 0</code> | Define the network (10.10.21.0/24) on which OSPF runs and associate the area ID (0). |
| <code>(config-router)#network 192.168.0.90/32 area 0</code> | Set the IP address of the loopback interface to 192.168.0.63/32. |
| <code>(config-router)#commit</code> | Commit the transaction. |

Establish a Trunk with CSPF Disabled


OcNOS, Constrained Shortest Path First (CSPF) calculation is enabled by default. Typically, CSPF is disabled when all of the participating nodes do not support the required traffic engineering extensions and LSPs are configured manually to use an explicit path. In this case, an LSP is established only along the path specified by the operator.



Note: This example is based on the assumption that a minimal configuration exists on all participating routers as described in [Enable Label Switching - Minimal Configuration \(page 644\)](#).

Topology

Figure 23. Basic Topology



PE1 - RSVP-TE

| | |
|---|---|
| <code>#configure terminal</code> | Enter configure mode. |
| <code>(config)#rsvp-trunk T1</code> | Create an RSVP trunk T1 and enter the Trunk mode. |
| <code>(config-trunk)#no primary cspf</code> | Specify no primary cspf since CSPF is enabled by default. |
| <code>(config-trunk)#to 192.168.0.90</code> | Specify the IPv4 egress (destination point) for the LSP. |
| <code>(config-trunk)#commit</code> | Commit the transaction. |

Establish a Trunk Using CSPF

The RSVP trunk can be configured using CSPF (Constraint-based Shortest Path First). In this case, the RSVP daemon (rsvpd) sends a request to the CSPF server to compute a path through the network to reach the destination. CSPF returns a hop-by-hop path called the Explicit Route to the RSVP daemon to be used in the Explicit Route Object (ERO). Each router along the path sends a `Path` message only to the nexthop specified in the ERO. In the OcNOS implementation, CSPF is enabled by default and if `no cspf` is not specified, the trunk is CSPF enabled automatically.



Note: This example is based on the assumption that a minimal configuration exists on all participating routers as described in [Enable Label Switching - Minimal Configuration \(page 644\)](#).

PE1 (RSVP Daemon)

| | |
|--------------------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#rsvp-trunk T1 | Create an RSVP trunk T1 and enter the Trunk mode. |
| (config-trunk)#to 192.168.0.90 | Specify the IPv4 egress (destination point) for the LSP. |
| (config-trunk)#commit | Commit the transaction. |

Mapping a Route to a Trunk


In the OcNOS implementation, a network can be mapped to a particular trunk using map-route configuration.



Note: This example is based on the assumption that a minimal configuration exists on all participating routers. For configuration details, refer to the “Establishing a Trunk - Minimal Configuration” section.

Topology

Figure 24. Topology for route mapping



PE1 - RSVP-TE

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router rsvp | Enable RSVP globally. |
| (config-router)#rsvp-path PE2 | Configure RSVP path. |
| (config-path)#rsvp-trunk T1 | Create an RSVP trunk T1 and enter the Trunk mode. |
| (config-trunk)#map-route 90.90.90.0/24 | Specify the destination prefix that needs to mapped to this trunk. |

| | |
|--------------------------------|--|
| (config-trunk)#to 192.168.0.90 | Specify the IPv4 egress (destination point) for the LSP. |
| (config-trunk)#commit | Commit the transaction. |

Establish a Trunk Using Explicitly-Defined Path

Explicit Route hops can be configured manually in the trunk configuration. In this case, the RSVP daemon uses the configured hops as Explicit Route Objects (ERO). It sets up the LSP using specified hops only.

An ERO is composed of IP addresses called hops. An ERO hop can be defined as loose or strict. A loose hop can be reached by any available route. A strict hop must be reached via a direct link and cannot be routed over any alternate routers in between. In this example, since PE3 is defined as loose hop, P can use P2 as an intermediate hop to reach PE2. However, if it was a strict hop, then P would have to use interface `xe1` to reach PE3 directly.




Notes:

- This example is based on the assumption that a minimal configuration exists on all participating routers as described in [Enable Label Switching - Minimal Configuration \(page 644\)](#).
- For more details on the hop configuration, refer to [A.B.C.D \(page 872\)](#).

Topology

Figure 25. Topology for Explicitly Defined Path



PE1 - RSVP-Path

| | |
|----------------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#rsvp-path P1 | Create an RSVP Path P1 and enter the Path mode. |
| (config-path)#10.10.23.60 strict | Configure this explicit route path as a <code>strict</code> hop. |
| (config-path)#10.10.21.90 loose | Configure this explicit route path as a <code>loose</code> hop. |
| (config-path)#exit | Exit Path mode. |
| #configure terminal | Enter configure mode. |
| (config)#rsvp-trunk T1 | Create an RSVP trunk T1 and enter the Trunk mode. |
| (config-trunk)#no primary cspf | Since CSPF is enabled by default, specify <code>no primary cspf</code> if CSPF is not required. |

| | |
|--------------------------------|--|
| (config-trunk)#primary path P1 | Configure trunk T1 to use the defined path. |
| (config-trunk)#to 192.168.0.90 | Specify the IPv4 egress (destination point) for the LSP. |
| (config-trunk)#commit | Commit the transaction. |

Validation

```
PE1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
```

```
Ingress RSVP:
To          From          Type    LSPName          State
Uptime    Rt  Style  Labelin  Labelout  DSType
192.168.0.90 192.168.0.63  PRI    T1-Primary      UP    00:00:16  1 1 SE  -
24320      DEFAULT
Total 1 displayed, Up 1, Down 0.
```

```
PE2#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to
ondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
```

```
Egress RSVP:
To          From          Type    LSPName          State Uptime    Rt  Style  Labelin  Labelout  DSType
192.168.0.90 192.168.0.63  PRI    T1-Primary      UP    00:00:33  1 1 SE  25600  -    ELSP_CON
Total 1 displayed, Up 1, Down 0.
PE1#show ip ospf neighbor
Total number of full neighbors: 1 OSPF process 100 VRF(default):
Neighbor ID Pri  State      Dead Time Address  Interface  Instance ID
192.168.0.60 1  Full/Backup 00:00:30  10.10.23.60 xe2        0
```

```
P1#show ip ospf neighbor
Total number of full neighbors: 2 OSPF process 100 VRF(default):
Neighbor ID Pri  State      Dead Time Address  Interface  Instance ID
192.168.0.90 1  Full/Backup 00:00:35  10.10.21.90 xe8        0
192.168.0.63 1  Full/DR    00:00:36  10.10.23.63 xe2        0
```

```
PE2#show ip ospf neighbor
Total number of full neighbors: 1 OSPF process 100 VRF(default):
Neighbor ID Pri  State      Dead Time Address  Interface  Instance ID
192.168.0.60 1  Full/DR    00:00:32  10.10.21.60 xe8        0
```

Add a Secondary LSP to the Trunk

Although the attributes of a Secondary LSP are independent of the Primary LSP, a Secondary LSP cannot be configured without first configuring a Primary LSP. In addition to information on how to configure a secondary LSP, this example illustrates how to define a non-default setup and the hold priority for an LSP. Setup and hold priorities are used to determine which LSP should be given a preference when competing for resources. Specifically, the setup priority of an un-established LSP is compared to the hold priorities of established LSPs, and the numerically lower one is given a preference. However, once the LSP is established, its setup priority is never used until it is pre-empted or reset and is being brought up again.

**Notes:**

- This example is based on the assumption that a minimal configuration exists on all participating routers as described in [Enable Label Switching - Minimal Configuration \(page 644\)](#).
- If user provides the RSVP path option for secondary, the primary path exclusion logic gets disabled. User needs to keep primary and secondary path mutually exclusive. Else, RSVP-Primary LSP and RSVP-Secondary LSP may select the same next hop, when RSVP is configured with "loose". Hence RSVP-Path first next-hop should be "strict".

PE1 - RSVP-TE

| | |
|---------------------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#rsvp-path myPath | Specify an RSVP path to be used. |
| (config-path)#10.10.23.60 strict | Configure this explicit route path as a <code>strict</code> hop. |
| (config-path)#exit | Exit Path mode. |
| (config)#rsvp-path myPath2 | Specify an RSVP path to be used. |
| (config-path)#10.10.23.60 loose | Configure this explicit route path as a <code>loose</code> hop. |
| (config-path)#exit | Exit Path mode. |
| (config)#rsvp-trunk T1 | Create an RSVP trunk <code>T1</code> and enter the Trunk mode. |
| (config-trunk)#no primary cspf | Since CSPF is enabled by default, specify <code>no secondary cspf</code> if CSPF is not required. |
| (config-trunk)#primary path myPath | Specify an RSVP path to be used. |
| (config-trunk)#no secondary cspf | Specify the <code>no secondary cspf</code> option for the Secondary LSP. |
| (config-trunk)#secondary path myPath2 | Specify an RSVP path to be used. |
| (config-trunk)#to 192.168.0.90 | Specify the IPv4 egress (destination point) for the LSP. |
| (config-trunk)#commit | Commit the transaction. |

Validation

This example shows the number of configured RSVP sessions in a router.

PE1

```
#show rsvp session count
Total configured: 1, Up 1, Down 0

Total ingress sessions: 1, Up 1, Down 0
Total transit sessions: 0, Up 0, Down 0
Total egress sessions: 0, Up 0, Down 0
```

PE2

```
#show rsvp session count
Total configured: 1, Up 1, Down 0

Total ingress sessions: 0, Up 0, Down 0
Total transit sessions: 0, Up 0, Down 0
Total egress sessions: 1, Up 1, Down 0
```

Add Multiple Secondary LSP to the trunk


RSVP Multiple Secondary feature tries to provide continuous protection when multiple failures happen. In majority scenarios, feature tries to provide seamless protection. This is a proprietary feature where user can configure multiple secondary sessions in a rsvp-trunk. Each secondary will be associated with a priority. Priority secondary sessions must be programmed with a predefined path. User can configure a maximum of five priority levels. Lowest priority number corresponds to highest priority. Highest priority session will be signaled to be programmed as secondary session. If highest priority session cannot come up, then next available secondary will be selected based on polling. During primary session fail-over, programmed secondary priority session will protect the primary and then goes for an MBB update to act as the primary session until primary comes up. Once the highest priority session comes up as acting primary session, next available secondary priority session will be programmed to signal and come up secondary. Re-optimization timer executed once in every 5 minutes to ensure the best priority session serves as secondary. Configuration updates on secondary priority configurations doesn't trigger MBB and session will be restarted. This example illustrates how to create SVI, enable IGP protocols and RSVP on SVI.



Note: Ensure that the VLAN is configured before creating SVI.

Topology

Figure 26. Topology for Multiple Secondary Protection

**Bridge Configuration**

```
bridge 1 protocol ieee vlan-bridge
no bridge 1 spanning-tree enable bridge-forward
```


VLAN creation

```

vlan database
vlan 2-7 bridge 1 state enable
vlan 501-506 bridge 1 state enable

```

PE1

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 47.47.47.47/32 secondary | Set the secondary IP address for the lo interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router rsvp | Enable RSVP globally. |
| (config-router)#exit | Exit RSVP mode. |
| (config)#interface vlan1.2 | Enter the interface mode. |
| (config-if)#ip address 10.10.23.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.3 | Enter the interface mode. |
| (config-if)#ip address 10.10.24.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.4 | Enter the interface mode. |
| (config-if)#ip address 10.10.25.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.5 | Enter the interface mode. |
| (config-if)#ip address 10.10.26.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |

| | |
|---|--|
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.6 | Enter the interface mode. |
| (config-if)#ip address 10.10.27.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.7 | Enter the interface mode. |
| (config-if)#ip address 10.10.28.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe2 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 2,501 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 501 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe3 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 3,502 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 502 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe4 | Enter the interface mode. |

| | |
|---|--|
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 4,503 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 503 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe5 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 5,504 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 504 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe6 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 6,505 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 505 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe7 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 7,506 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 506 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (conf)#rsvp-path p1-r1-r3 mpls | Create RSVP path |

| | |
|--|---|
| (conf-path)#10.10.23.2 strict | Configure nexthop |
| (conf-path)#10.10.21.2 strict | Configure nexthop |
| (conf)#rsvp-path sp1-r1-r3 mpls | Create RSVP path |
| (conf-path)#10.10.24.2 strict | Configure nexthop |
| (conf-path)#10.10.22.2 strict | Configure nexthop |
| (conf)#rsvp-path sp2-r1-r3 mpls | Create RSVP path |
| (conf-path)#10.10.25.2 strict | Configure nexthop |
| (conf-path)#10.10.29.2 strict | Configure nexthop |
| (conf)#rsvp-path sp3-r1-r3 mpls | Create RSVP path |
| (conf-path)#10.10.26.2 strict | Configure nexthop |
| (conf-path)#10.10.30.2 strict | Configure nexthop |
| (conf)#rsvp-path sp4-r1-r3 mpls | Create RSVP path |
| (conf-path)#10.10.27.2 strict | Configure nexthop |
| (conf-path)#10.10.31.2 strict | Configure nexthop |
| (conf)#rsvp-path sp5-r1-r3 mpls | Create RSVP path |
| (conf-path)#10.10.28.2 strict | Configure nexthop |
| (conf-path)#10.10.32.2 strict | Configure nexthop |
| (conf)#rsvp-trunk 47-49-test ipv4 | Create a RSVP trunk link |
| (conf-trunk)#primary path p1-r1-r3 | Configure primary path for trunk link |
| (conf-trunk)#secondary-priority 1 path sp1-r1-r3 | Configure secondary link for trunk link |
| (conf-trunk)#secondary-priority 2 path sp2-r1-r3 | Configure secondary link for trunk link |
| (conf-trunk)#secondary-priority 3 path sp3-r1-r3 | Configure secondary link for trunk link |
| (conf-trunk)#secondary-priority 4 path sp4-r1-r3 | Configure secondary link for trunk link |
| (conf-trunk)#secondary-priority 5 path sp5-r1-r3 | Configure secondary link for trunk link |
| (conf-trunk)#to 49.49.49.49 | Configure remote node for the LSP |
| (config-trunk)#commit | Commit the transaction. |

OSPF configurations

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the Routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 47.47.47.47 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 10.10.23.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.24.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |

| | |
|---|---|
| (config-router)#network 10.10.25.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.26.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.27.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.28.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 47.47.47.47/32 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router) #commit | Commit the transaction. |

Validation

This example shows the number of configured RSVP sessions in a router.

PE1

```
#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
```

Ingress RSVP:

| To | From | Type | LSPName | State |
|-------------|-------------|-------|---------------------------------|-------------------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 49.49.49.49 | 47.47.47.47 | PRI | 47-49-test-Primary | UP |
| 24961 | DEFAULT | | | 00:32:35 1 1 SE - |
| 49.49.49.49 | 47.47.47.47 | SEC | 47-49-test-Secondary-Priority-1 | UP |
| 24962 | DEFAULT | | | 00:32:35 1 1 SE - |

Total 2 displayed, Up 2, Down 0.

Egress RSVP:

| To | From | Type | LSPName | State |
|-------------|-------------|-------|---------------------------------|--------------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 47.47.47.47 | 49.49.49.49 | PRI | 49-47-test-Primary | UP |
| SE | 24964 | - | ELSP_CON | 00:32:53 1 1 |
| 47.47.47.47 | 49.49.49.49 | PRI | 49-47-test-Secondary-Priority-1 | UP |
| SE | 24962 | - | ELSP_CON | 00:32:47 1 1 |

Total 2 displayed, Up 2, Down 0.

#show rsvp trunk multi-sec-detail

```
Ingress (Secondary-Priority1)
49.49.49.49
  From: 47.47.47.47, LSPstate: Up, LSPname: 47-49-test-Secondary-Priority-1
  Ingress FSM state: Operational
  Establishment Time: 0s 253ms
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
  IGP-Shortcut: Disabled, LSP metric: 3
  LSP Protection: None
  Label in: -, Label out: 24962,
  Tspecc rate: 0, Fspecc rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2219, Ext-Tunnel Id: 47.47.47.47
  Downstream: 47.46.3.2, vlan1.1003
  Path refresh: 30 seconds (RR enabled) (due in 27970 seconds)
  Resv lifetime: 157 seconds (due in 138 seconds)
```

```

Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: SP1-47-49 (in use)
Configured Explicit Route Detail :
  47.46.3.2/32 strict
  46.45.9.2/32 strict
  45.49.24.2/32 strict
Session Explicit Route Detail :
  47.46.3.2/32 strict
  46.45.9.2/32 strict
  45.49.24.2/32 strict
Record route:
-----
IP Address      Label
-----
<self>
47.46.3.2
46.45.9.2
45.49.24.2
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 9216
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Ingress (Secondary-Priority2)
49.49.49.49
  From: 47.47.47.47, LSPstate: Dn, LSPname: 47-49-test-Secondary-Priority-2
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: NA
  IGP-Shortcut: Disabled, LSP metric: 3
  LSP Protection: None
  Label in: -, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2223, Ext-Tunnel Id: 47.47.47.47
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Ingress (Secondary-Priority3)
49.49.49.49
  From: 47.47.47.47, LSPstate: Dn, LSPname: 47-49-test-Secondary-Priority-3
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: NA
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: -, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2219, Ext-Tunnel Id: 47.47.47.47
  Last Recorded Error Code: Routing Problem (24)
  Last Recorded Error Value: No route available toward destination (5)
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Ingress (Secondary-Priority4)
49.49.49.49
  From: 47.47.47.47, LSPstate: Dn, LSPname: 47-49-test-Secondary-Priority-4
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds

```

```

LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: NA
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: -, Label out: -,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5001, LSP Id: 2219, Ext-Tunnel Id: 47.47.47.47
Last Recorded Error Code: Routing Problem (24)
Last Recorded Error Value: No route available toward destination (5)
Node where Last Recorded Error originated: None
Trunk Type: mpls
Ingress (Secondary-Priority5)
49.49.49.49
  From: 47.47.47.47, LSPstate: Dn, LSPname: 47-49-test-Secondary-Priority-5
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: NA
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: -, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2219, Ext-Tunnel Id: 47.47.47.47
  Last Recorded Error Code: Routing Problem (24)
  Last Recorded Error Value: No route available toward destination (5)
  Node where Last Recorded Error originated: None
  Trunk Type: mpls

```

P Bridge Configuration

```

bridge 1 protocol ieee vlan-bridge
no bridge 1 spanning-tree enable bridge-forward

```

VLAN creation (Peer configuration for PE1)

```

vlan database
vlan 2-7 bridge 1 state enable
vlan 507-512 bridge 1 state enable

```

VLAN creation (Peer configuration for PE2)

```

vlan database
vlan 9-14 bridge 1 state enable
vlan 513-518 bridge 1 state enable

```

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 46.46.46.46/32 secondary | Set the secondary IP address for the lo interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router rsvp | Enable RSVP globally. |
| (config-router)#exit | Exit RSVP mode. |
| (config)#interface vlan1.2 | Enter the interface mode. |
| (config-if)#ip address 10.10.23.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |

| | |
|--|--|
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.3 | Enter the interface mode. |
| (config-if)#ip address 10.10.24.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.4 | Enter the interface mode. |
| (config-if)#ip address 10.10.25.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.5 | Enter the interface mode. |
| (config-if)#ip address 10.10.26.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.6 | Enter the interface mode. |
| (config-if)#ip address 10.10.27.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.7 | Enter the interface mode. |
| (config-if)#ip address 10.10.28.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |

| | |
|--|--|
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.9 | Enter the interface mode. |
| (config-if)#ip address 10.10.21.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.10 | Enter the interface mode. |
| (config-if)#ip address 10.10.22.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.11 | Enter the interface mode. |
| (config-if)#ip address 10.10.29.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.12 | Enter the interface mode. |
| (config-if)#ip address 10.10.30.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.13 | Enter the interface mode. |
| (config-if)#ip address 10.10.31.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |

| | |
|---|--|
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.14 | Enter the interface mode. |
| (config-if)#ip address 10.10.32.1/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe2 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 2,507 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 507 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe3 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 3,508 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 508 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe4 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 4,509 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 509 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |

| | |
|---|--|
| (config)#interface xe5 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 5,510 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 510 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe6 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 6,511 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 511 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe7 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 7,512 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 512 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe9 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 9,513 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 513 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |

| | |
|--|--|
| (config)#interface xe10 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 10,514 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 514 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe11 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 11,515 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 515 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe12 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 12,516 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 516 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe13 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 13,517 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 517 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |

| | |
|--|--|
| (config)#interface xe14 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 14,518 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 518 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#commit | Commit the transaction. |

OSPF configurations

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the Routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 46.46.46.46 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 10.10.21.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.22.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.23.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.24.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.25.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.26.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.27.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.28.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.29.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.30.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.31.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.32.0/24 area 0 | Define the network on which OSPF runs and |

| | |
|---|---|
| | associate the area ID |
| (config-router)#network 46.46.46.46/32 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router) #commit | Commit the transaction. |

PE2 Bridge Configuration

```
bridge 1 protocol ieee vlan-bridge
no bridge 1 spanning-tree enable bridge-forward
```

PE2 VLAN Creation

```
vlan database
vlan 9-14 bridge 1 state enable
vlan 519-524 bridge 1 state enable
```

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 49.49.49.49/32 secondary | Set the secondary IP address for the lo interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router rsvp | Enable RSVP globally. |
| (config-router)#exit | Exit RSVP mode. |
| (config)#interface vlan1.9 | Enter the interface mode. |
| (config-if)#ip address 10.10.21.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.10 | Enter the interface mode. |
| (config-if)#ip address 10.10.22.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.11 | Enter the interface mode. |
| (config-if)#ip address 10.10.29.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |

| | |
|---|--|
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.12 | Enter the interface mode. |
| (config-if)#ip address 10.10.30.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.13 | Enter the interface mode. |
| (config-if)#ip address 10.10.31.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface vlan1.14 | Enter the interface mode. |
| (config-if)#ip address 10.10.32.2/24 | Configure the IP Address |
| (config-if)#mtu 1600 | Configure MTU size. |
| (config-if)#label-switching | Enable MPLS. |
| (config-if)#ip ospf network point-to-point | Enable OSPF point-to-point network type. |
| (config-if)#enable-rsvp | Enable RSVP at the interface level. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe9 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 9,519 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 519 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe10 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |

| | |
|--|--|
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 10,520 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 520 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe11 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 11,521 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 521 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe12 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 12,522 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 522 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe13 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 13,523 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 523 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe14 | Enter the interface mode. |
| (config-if)#switchport | Configure Switchport |
| (config-if)#bridge-group 1 | Assign a Bridge ID to the port channel |

| | |
|--|---|
| (config-if)#switchport mode trunk | Configure trunk |
| (config-if)#switchport trunk allowed vlan add 14,524 | Configure allowed VLANs |
| (config-if)#switchport trunk native vlan 524 | Configure native VLAN. |
| (config-if)#load-interval 30 | Set load interval |
| (config-if)#mtu 9192 | Configure the MTU Size. |
| (config-if)#exit | Exit the interface mode. |
| (conf)#rsvp-path sp1-r3-r1 mpls | Create RSVP path |
| (conf-path)#10.10.21.1 strict | Configure nexthop |
| (conf-path)#10.10.23.1 strict | Configure nexthop |
| (conf)#rsvp-path sp2-r3-r1 mpls | Create RSVP path |
| (conf-path)#10.10.22.1 strict | Configure nexthop |
| (conf-path)#10.10.24.1 strict | Configure nexthop |
| (conf)#rsvp-path sp2-r3-r1 mpls | Create RSVP path |
| (conf-path)#10.10.29.2 strict | Configure nexthop |
| (conf-path)#10.10.25.1 strict | Configure nexthop |
| (conf)#rsvp-path sp3-r3-r1 mpls | Create RSVP path |
| (conf-path)#10.10.30.1 strict | Configure nexthop |
| (conf-path)#10.10.26.1 strict | Configure nexthop |
| (conf)#rsvp-path sp4-r3-r1 mpls | Create RSVP path |
| (conf-path)#10.10.31.1 strict | Configure nexthop |
| (conf-path)#10.10.27.1 strict | Configure nexthop |
| (conf)#rsvp-path sp5-r3-r1 mpls | Create RSVP path |
| (conf-path)#10.10.32.1 strict | Configure nexthop |
| (conf-path)#10.10.28.1 strict | Configure nexthop |
| (conf)#rsvp-trunk 49-47-test ipv4 | Create a RSVP trunk link |
| (conf-trunk)#primary path p1-r3-r1 | Configure primary path for trunk link |
| (conf-trunk)#secondary-priority 1 path sp1-r3-r1 | Configure secondary link for trunk link |
| (conf-trunk)#secondary-priority 2 path sp2-r3-r1 | Configure secondary link for trunk link |
| (conf-trunk)#secondary-priority 3 path sp3-r3-r1 | Configure secondary link for trunk link |
| (conf-trunk)#secondary-priority 4 path sp4-r3-r1 | Configure secondary link for trunk link |
| (conf-trunk)#secondary-priority 5 path sp5-r3-r1 | Configure secondary link for trunk link |
| (conf-trunk)#to 47.47.47.47 | Configure remote node for the LSP |
| (conf-trunk)#commit | Commit the transaction. |

OSPF configurations

| | |
|--------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the Routing process and specify the |

| | |
|---|---|
| | Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 49.49.49.49 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 10.10.21.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.22.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.29.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.30.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.31.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 10.10.32.0/24 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#network 49.49.49.49/32 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router) #commit | Commit the transaction. |

Validation

This example shows the number of configured RSVP sessions in a router.

PE2

```
#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Type    LSPName          State
Uptime    Rt  Style  Labelin  Labelout  DSType
47.47.47.47 49.49.49.49 PRI     49-47-test-Primary UP    00:34:57  1 1 SE  -
24970      DEFAULT
47.47.47.47 49.49.49.49 SEC     49-47-test-Secondary-Priority-1 UP    00:34:56  1 1 SE  -
24968      DEFAULT
Total 2 displayed, Up 2, Down 0.

Egress RSVP:
To          From          Type    LSPName          State
Uptime    Rt  Style  Labelin  Labelout  DSType
49.49.49.49 47.47.47.47 PRI     47-49-test-Primary UP    00:34:45  1 1
SE    31364  -      ELSP_CON
49.49.49.49 47.47.47.47 PRI     47-49-test-Secondary-Priority-1 UP    00:34:44  1 1
SE    31360  -      ELSP_CON
Total 2 displayed, Up 2, Down 0.

#show rsvp trunk multi-sec-detail
Ingress (Secondary-Priority1)
47.47.47.47
  From: 49.49.49.49, LSPstate: Up, LSPname: 49-47-test-Secondary-Priority-1
  Ingress FSM state: Operational
```

```

Establishment Time: 1s 71ms
Setup priority: 7, Hold priority: 0
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
IGP-Shortcut: Disabled, LSP metric: 3
LSP Protection: None
Label in: -, Label out: 24968,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5001, LSP Id: 2214, Ext-Tunnel Id: 49.49.49.49
Downstream: 45.49.24.1, vlan1.1024
Path refresh: 30 seconds (RR enabled) (due in 27829 seconds)
Resv lifetime: 157 seconds (due in 145 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: SP1-49-47 (in use)
Configured Explicit Route Detail :
  45.49.24.1/32 strict
  46.45.9.1/32 strict
  47.46.3.1/32 strict
Session Explicit Route Detail :
  45.49.24.1/32 strict
  46.45.9.1/32 strict
  47.46.3.1/32 strict
Record route:
-----
IP Address      Label
-----
<self>
45.49.24.1
46.45.9.1
47.46.3.1
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 9216
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Ingress (Secondary-Priority2)
47.47.47.47
  From: 49.49.49.49, LSPstate: Dn, LSPname: 49-47-test-Secondary-Priority-2
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: NA
  IGP-Shortcut: Disabled, LSP metric: 3
  LSP Protection: None
  Label in: -, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2215, Ext-Tunnel Id: 49.49.49.49
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Ingress (Secondary-Priority3)
47.47.47.47
  From: 49.49.49.49, LSPstate: Dn, LSPname: 49-47-test-Secondary-Priority-3
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: NA
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: -, Label out: -,

```

```

Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5001, LSP Id: 2213, Ext-Tunnel Id: 49.49.49.49
Last Recorded Error Code: Routing Problem (24)
Last Recorded Error Value: No route available toward destination (5)
Node where Last Recorded Error originated: None
Trunk Type: mpls
Ingress (Secondary-Priority4)
47.47.47.47
  From: 49.49.49.49, LSPstate: Dn, LSPname: 49-47-test-Secondary-Priority-4
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: NA
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: -, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2213, Ext-Tunnel Id: 49.49.49.49
  Last Recorded Error Code: Routing Problem (24)
  Last Recorded Error Value: No route available toward destination (5)
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Ingress (Secondary-Priority5)
47.47.47.47
  From: 49.49.49.49, LSPstate: Dn, LSPname: 49-47-test-Secondary-Priority-5
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: NA
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: -, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2213, Ext-Tunnel Id: 49.49.49.49
  Last Recorded Error Code: Routing Problem (24)
  Last Recorded Error Value: No route available toward destination (5)
  Node where Last Recorded Error originated: None
  Trunk Type: mpls

```

Add Administrative Group Constraints to an LSP

To add administrative group constraints (also known as color constraints) to an LSP:

- Configure support for required administrative groups in NSM on all participating routers
- Configure required administrative groups on all participating interfaces

The configuration in this example forces the primary LSP to be setup through links that belong either to administrative group A or C. A link that does not belong to either of these administrative groups will not be used for setting up the LSP.



Note: This example is based on the assumption that a minimal configuration exists on all participating routers as described in [Enable Label Switching - Minimal Configuration \(page 644\)](#).

PE1 - NSM

```
#configure terminal
```

Enter configure mode.

| | |
|---------------------------|--|
| (config)#admin-group A 0 | Add new administrative groups, specify their names and assign bit values to them. |
| (config)#admin-group B 1 | |
| (config)#admin-group C 2 | |
| (config)#admin-group D 3 | |
| (config)#admin-group E 4 | |
| (config)#interface xe0 | Enter interface mode. |
| (config-if)#admin-group A | Add administrative groups to the links. When used in the interface mode, this command adds a link between an interface and a group. The name is the name of the group previously configured. You can have multiple groups per interface. |
| (config-if)#admin-group B | |
| (config-if)#admin-group C | |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#admin-group E | Add administrative groups to the links. When used in the interface mode, this command adds a link between an interface and a group. The name is the name of the group previously configured. You can have multiple groups per interface. |
| (config-if)#admin-group D | |
| (config-if)#commit | Commit the transaction. |

P - NSM

| | |
|---------------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#admin-group A 0 | Add new administrative groups and specify their names and assign bit values to them. |
| (config)#admin-group C 2 | |
| (config)#interface xe2 | Enter interface mode |
| (config-if)#admin-group A | Add administrative groups to the links. When used in the interface mode, this command adds a link between an interface and a group. The name is the name of the group previously configured. You can have multiple groups per interface. |
| (config-if)#admin-group C | |
| (config-if)#commit | Commit the transaction. |

PE1 - RSVP-TE

| | |
|--------------------------------|--|
| (config)#router rsdp | Enter the router mode for RSVP. |
| (config-router)#rsdp-path P1 | Configure RSVP path. |
| #configure terminal | Enter configure mode. |
| (config)#rsdp-trunk T1 | Create an RSVP trunk T1 and enter the Trunk mode. |
| (config-trunk)#no primary cspf | Since CSPF is enabled by default, specify no primary cspf if CSPF is not required. |

| | |
|---|--|
| <code>(config-trunk)#primary path P1</code> | Specify an RSVP primary path to be used. |
| <code>(config-trunk)#no primary cspf</code> | Specify the <code>no primary cspf</code> option for the LSP. |
| <code>(config-trunk)#primary include-any A</code> | Set up the LSP with admin group constraint A. |
| <code>(config-trunk)#primary include-any C</code> | Set up the LSP with admin group constraint C. |
| <code>(config-trunk)#to 192.168.0.90</code> | Specify the IPv4 egress (destination point) for the LSP. |
| <code>(config-trunk)#commit</code> | Commit the transaction. |

Configure Global Parameters

Some common parameters can be configured in the Router mode on the RSVP-TE daemon. These parameters are global and affect all LSPs. In the following example the interval between two consecutive hello messages is set. The neighbor is defined by the `neighbor` command. Hello exchanges are enabled only between explicitly configured neighbors (configure this router as a neighbor on P (IP address 10.10.23.60)).



Note: This example is based on the assumption that a minimal configuration exists on all participating routers as described in [Enable Label Switching - Minimal Configuration \(page 644\)](#).

PE1 - RSVP-TE

| | |
|---|--|
| <code>#configure terminal</code> | Enter configure mode. |
| <code>(config)#router rsvp</code> | Enter the router mode for RSVP. |
| <code>(config-router)#hello-interval 10</code> | Set the <code>hello-interval</code> (in seconds) between hello packets. |
| <code>(config-router)#hello-timeout 35</code> | Set the <code>hello-timeout</code> value. If an LSR has not received a Hello message from a peer within this period, all sessions shared with this peer are reset. |
| <code>(config-router)#neighbor 10.10.23.60</code> | Explicitly specify the neighbor to exchange Hello messages with. |
| <code>(config-router)#commit</code> | Commit the transaction. |

P - RSVP-TE

| | |
|---|--|
| <code>#configure terminal</code> | Enter configure mode. |
| <code>(config)#router rsvp</code> | Enter the router mode for RSVP. |
| <code>(config-router)#hello-interval 10</code> | Set the <code>hello-interval</code> (in seconds) between hello packets. |
| <code>(config-router)#hello-timeout 35</code> | Set the <code>hello-timeout</code> value. If an LSR has not received a Hello message from a peer within this period, all sessions shared with this peer are reset. |
| <code>(config-router)#neighbor 10.10.23.63</code> | Explicitly specify the neighbor to exchange Hello |

| | |
|------------------------|-------------------------|
| | messages with. |
| (config-router)#commit | Commit the transaction. |

MPLS RSVP PING and TRACEROUTE

This example shows MPLS ping and trace route for RSVP

```
#ping mpls rsvp tunnel-name to_30 detail
Sending 5 MPLS Echos to to_30 , timeout is 5 seconds
Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed
Type 'Ctrl+C' to abort
! seq_num = 1 30.30.30.30 0.28 ms
! seq_num = 2 30.30.30.30 0.24 ms
! seq_num = 3 30.30.30.30 0.22 ms
! seq_num = 4 30.30.30.30 0.22 ms
! seq_num = 5 30.30.30.30 0.22 ms

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.22/0.25/0.28

RTR-29#trace mpls rsvp tunnel-name to_30 detail
Tracing MPLS Label Switched Path to to_30 , timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

  0 21.21.21.29 [Labels: 24320]
R 1 43.43.43.43 [Labels: 24320] 123.22 ms
R 2 42.42.42.42 [Labels: 24960] 1.60 ms
! 3 30.30.30.30 1.62 ms
```

MPLS RSVP Entropy Label Capabilities

To share the load across multiple members of a LAG port in the core of an MPLS network we can use entropy labels. An Entropy Label is always preceded by an Entropy level indicator which is a special Label with value seven, and indicates the next label present is an Entropy label. The trade off is the MPLS stack depth increases by two and it reduces overhead on transit routers.



Note: Load balancing is enabled by default for all the parameters. If you enable load balancing manually, then all the parameters enabled by default are reset and you need to enable the parameters based on which traffic should be load balanced.

Entropy labels will only be added when the remote edge node advertises its capability for Entropy.

The examples below show how entropy can be enabled on a provider edge node as per the setup we need to enable on RTR29 and RTR30 to have entropy enabled in both directions:

```
(config)#router rsvp
(config-router)#entropy-label-capability
```

This enables ELC signaling for RSVP.

For validation, use:

```
#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, B - BGP FTN, K - CLI FTN, t - tunnel
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code      FEC          FTN-ID Tunnel-id    Pri  LSP-Type    Out- Label  ELC  Out-Intf  Nexthop
R(t)> 2  9.29.29.29/32  1      5001      Yes  LSP_DEFAULT  24322    yes
xe2      1.41.41.31
R(t)> 29.29.29.29/32  2      5001      No   LSP_DEFAULT  4322     yes
xe1      69.69.69.42

#show rsvp session ingress detail
Ingress (Primary)
41.41.41.31
From: 29.29.29.29, LSPstate: Up, LSPname: t1-Primary
Ingress FSM state: Operational
Establishment Time: 0s 3ms
Setup priority: 7, Hold priority: 0
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
IGP-Shortcut: Disabled, LSP metric: 1
LSP Protection: None
Label in: -, Label out: 24320, ELC
```

MPLS RSVP LSP Re-optimization

Follow these steps to configure RSVP LSP Re-optimization.

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#rsvp-trunk T1 | Create an RSVP trunk T1 and enter the Trunk mode. |
| (config-trunk)#reoptimize | Enable re-optimization of the session. |
| #configure terminal | Enter configure mode. |
| (config)#router rsvp | Enter RSVP mode |
| (config-router)#lsp-reoptimization-timer 5 | Sets the re-optimization timer for the session. |

Follow these steps to force the LSP to be re-optimized.

| | |
|---|--------------------------------|
| #rsvp-session T1-Primary force-reoptimize | Re-optimize the LSP forcefully |
|---|--------------------------------|

RSVP-TE Facility Backup (Facility Bypass)

RSVP supports multiple path protection mechanisms and facility backup is one of them. With facility backup protection, N number of LSPs sharing the common path can be protected using one bypass tunnel which leads to resource utilization.




Notes:

- Do not configure a facility backup trunk with the same transit node as that of the primary trunk.
- IP BFD must be configured at 3.3ms to achieve 50ms switchover for RSVP FRR.

Topology

As shown in [Figure 27](#), we have four routers R1, R2, R3, and R4.

Figure 27. RSVP facility backup



Configuration

PE1

| | |
|---|--|
| #configure terminal | Enter configuration mode |
| (config)#interface lo | Specify interface loopback for configuration |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure ip address of loopback |
| (config-if)#exit | Exit interface configuration mode |
| (config)#interface xel | Specify interface xe1 for configuration |

| | |
|--|---|
| (config-if)#ip address 10.1.2.12/24 | Configure ip address of interface |
| (config)#exit | Exit interface configuration mode |
| (config-if)#int xe5 | Specify interface xe1 for configuration |
| (config-if)#ip address 10.1.4.14/24 | Configure ip address of loopback |
| (config-if)#exit | Exit configuration mode |
| (config)#router ospf 1 | Configure the router OSPF with process id |
| (config-router)#router-id 1.1.1.1 | Configure OSPF router-id |
| (config-router)#network 1.1.1.1/32 area 1 | Define the network of the interface with area 0 |
| (config-router)#network 10.1.2.0/24 area 1 | Define the network of the interface with area 0 |
| (config-router)#network 10.1.4.0/24 area 1 | Define the network of the interface with area 0 |
| (config-router)#exit | Exit the configure mode |
| (config)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval |
| (config)#exit | Exit the configure mode |
| (config)#router ospf 1 | Enter router OSPF mode with process id |
| (config-router)#bfd all-interfaces | Enable the OSPF enabled interfaces with bfd |
| (config-router)#exit | Exit the router mode |
| (config)#router rsvp | Enter router RSVP |
| (config-router)#exit | Exit the router configuration mode |
| (config)#interface xe1 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP |
| (config-if)#label-switching | Enable label-switching |
| (config-if)#exit | Exit the interface configuration mode |
| (config)#interface xe5 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP |
| (config-if)#label-switching | Enable label-switching |
| (config-if)#commit | Commit the transaction. |

P1

| | |
|---|--|
| (config)#interface lo | Specify the interface (lo) |
| (config-if)#ip address 2.2.2.2/32 secondary | Enter the loopback ip address as secondary |
| (config-if)#exit | Exit the interface configure mode |
| (config-if)#int xe1 | Specify the interface(xe1) |
| (config-if)#ip address 10.1.2.21/24 | Configure the IP address for the interface |
| (config-if)#exit | Exit the interface mode |
| (config-if)#int xe10 | Specify the interface(xe1) |
| (config-if)#ip address 10.2.3.23/24 | Configure the IP address for the interface |

| | |
|--|---|
| (config-if)#exit | Exit the configuration mode |
| (config)#router ospf 1 | Configure OSPF router-id |
| (config-router)#router-id 2.2.2.2 | Configure the router id |
| (config-router)#network 2.2.2.2/32 area 1 | Define the network of the interface with area 0 |
| (config-router)#network 10.1.2.0/24 area 1 | Define the network of the interface with area 0 |
| (config-router)#network 10.2.3.0/24 area 1 | Define the network of the interface with area 0 |
| (config-router)#exit | Exit the configure mode |
| (config)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval |
| (config)#exit | Exit the configure mode |
| (config)#router ospf 1 | Enter router OSPF mode with process id |
| (config-router)#bfd all-interfaces | Enable the OSPF enabled interfaces with bfd |
| (config-router)#exit | Exit the router mode |
| (config)#router rsvp | Enter router RSVP |
| (config-router)#exit | Exit the router configuration mode |
| (config)#interface xe1 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP |
| (config-if)#label-switching | Enable label-switching |
| (config-if)#exit | Exit the interface configuration mode |
| (config)#interface xe10 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP |
| (config-if)#label-switching | Enable label-switching |
| (config-if)#commit | Commit the transaction. |

P2

| | |
|---|---|
| (config)#interface lo | Specify the interface (lo) |
| (config-if)#ip address 3.3.3.3/32 secondary | Enter the loopback ip address as secondary |
| (config-if)#exit | Exit the interface configuration mode |
| (config-if)#int xe10 | Specify the interface(xe1) |
| (config-if)#ip address 10.2.3.32/24 | Configure the IP address for the interface |
| (config-if)#exit | Exit the interface mode |
| (config-if)#interface xe5 | Specify the interface(xe1) |
| (config-if)#ip address 10.3.4.34/24 | Configure the IP address for the interface |
| (config-if)#exit | Exit the configuration mode |
| (config)#router ospf 1 | Configure OSPF router-id |
| (config-router)#router-id 3.3.3.3 | Configure the router id |
| (config-router)#network 3.3.3.3/32 area 1 | Define the network of the interface with area 0 |

| | |
|--|---|
| (config-router)#network 10.3.4.0/24 area 1 | Define the network of the interface with area 0 |
| (config-router)#network 10.2.3.0/24 area 1 | Define the network of the interface with area 0 |
| (config-router)#exit | Exit the configure mode |
| (config)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval |
| (config)#exit | Exit the configure mode |
| (config)#router ospf 1 | Enter router OSPF mode with process id |
| (config-router)#bfd all-interfaces | Enable the OSPF enabled interfaces with bfd |
| (config-router)#exit | Exit the router mode |
| (config)#router rsvp | Enter router RSVP |
| (config-router)#exit | Exit the router configuration mode |
| (config)#interface xe10 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP |
| (config-if)#label-switching | Enable label-switching |
| (config-if)#exit | Exit the interface configuration mode |
| (config)#interface xe5 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP |
| (config-if)#label-switching | Enable label-switching |
| (config-if)#commit | Commit the transaction. |

PE2

| | |
|--|---|
| (config)#interface lo | Specify the interface (lo) |
| (config-if)#ip address 4.4.4.4/32 secondary | Enter the loopback IP address as secondary |
| (config-if)#exit | Exit the interface configuration mode |
| (config-if)#interface xe2 | Specify the interface(xe1) |
| (config-if)#ip address 10.1.4.41/24 | Configure the ip address for the interface |
| (config-if)#exit | Exit the interface mode |
| (config-if)#int xe5 | Specify the interface(xe1) |
| (config-if)#ip address 10.3.4.43/24 | Configure the ip address for the interface |
| (config-if)#exit | Exit the configuration mode |
| (config)#router ospf 1 | Configure ospf router-id |
| (config-router)#router-id 4.4.4.4 | Configure the router id |
| (config-router)#network 4.4.4.4/32 area 1 | Define the network of the interface with area 0 |
| (config-router)#network 10.1.4.0/24 area 1 | Define the network of the interface with area 0 |
| (config-router)#network 10.3.4.0/24 area 1 | Define the network of the interface with area 0 |
| (config-router)#exit | Exit the configure mode |
| (config)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval |

| | |
|------------------------------------|---|
| (config)#exit | Exit the configuration mode |
| (config)#router ospf 1 | Exit the router OSPF mode with process id |
| (config-router)#bfd all-interfaces | Enable the OSPF enabled interfaces with bfd |
| (config-router)#exit | Exit the router mode |
| (config)#router rsvp | Enter router RSVP |
| (config-router)#exit | Exit the router configuration mode |
| (config)#interface xe1 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP |
| (config-if)#label-switching | Enable label-switching |
| (config-if)#exit | Exit the interface configuration mode |
| (config)#interface xe5 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP |
| (config-if)#label-switching | Enable label-switching |
| (config-if)#commit | Commit the transaction. |

RSVP Path on PE1

| | |
|---|--|
| (config)#rsvp-path primary_1 | Enter the rsvp-path configuration mode with name |
| (config-path)#10.1.2.21 strict | Specify the first next-hop ip address |
| (config-path)#10.2.3.32 strict | Specify the second next-hop ip address |
| (config-path)#exit | Exit the rsvp-path configuration mode |
| #configure terminal | Enter the configuration mode |
| (config)#rsvp-path bypass_1 | Enter the rsvp-path configuration mode with name |
| (config-path)#10.1.4.41 strict | Specify the first next-hop ip address |
| (config-path)#10.3.4.34 strict | Specify the second next-hop ip address |
| (config-path)#exit | Exit the rsvp-path configuration mode |
| #configure terminal | Enter the configuration mode |
| (config)#rsvp-trunk R1-R3 | Enter the rsvp trunk to be created with name |
| (config-trunk)#primary path primary_1 | Configure primary path for the trunk |
| (config-trunk)#to 3.3.3.3 | Enter the destination ip |
| (config-trunk)#primary fast-reroute protection facility | Configure facility backup protection for the trunk |
| (config-trunk)#exit | Exit the configuration mode |
| (config)#rsvp-bypass B1-B8 | Enter the rsvp bypass to be created with name |
| (config-trunk)#path bypass_1 | Configure primary path for the trunk |
| (config-trunk)#to 3.3.3.3 | Enter the destination IP |
| (config-if)#commit | Commit the transaction. |

Validation

OSPF Neighborhood

PE1

```
#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address      Interface     Instance ID
2.2.2.2        1     Full/Backup     00:00:38    10.1.2.21    xe1           0
4.4.4.4        1     Full/DR         00:00:33    10.1.4.41    xe5           0
```

P1

```
#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address      Interface     Instance ID
1.1.1.1        1     Full/DR         00:00:35    10.1.2.12    xe1           0
3.3.3.3        1     Full/Backup     00:00:34    10.2.3.32    xe10          0
```

P2

```
#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address      Interface     Instance ID
2.2.2.2        1     Full/DR         00:00:37    10.2.3.23    xe10          0
4.4.4.4        1     Full/Backup     00:00:39    10.3.4.43    xe5           0
```

PE2

```
#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address      Interface     Instance ID
1.1.1.1        1     Full/Backup     00:00:38    10.1.4.14    xe2           0
3.3.3.3        1     Full/DR         00:00:36    10.3.4.34    xe5           0
```

RSVP Session

PE1

```
#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes

Ingress RSVP:
To      From      Type LSPName      State
Uptime  Rt  Style Labelin Labelout DStype
3.3.3.3 24321 1.1.1.1 PRI R1-R3-Primary UP 00:54:48 1 1 SE -
3.3.3.3 1.1.1.1 DEFAULT BPS B1-B4-Bypass UP 01:08:32 1 1 SE -
```

```

24321    DEFAULT
Total 2 displayed, Up 2, Down 0.
#show rsvp bypass
Ingress RSVP:
To          From          LSPName          State
Uptime    Rt   Style  Labelin  Labelout  DStype
3.3.3.3    24321  DEFAULT  B1-B4-Bypass  UP      01:09:17  1 1 SE      -

#show rsvp bypass protected-lsp-list
Bypass trunk: B1-B4
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id    Lsp Id    Lsp Name          Role    Ext_tnl_
id      Ingress  Egress
5001      2202      R1-R3-
Primary    Ingress   1.1.1.1          1.1.1.1  3.3.3.3
Total LSP protected : 1
Bandwidth in use : 0

#show rsvp bypass B1-B4 protected-lsp-list
Bypass trunk: B1-B4
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id    Lsp Id    Lsp Name          Role    Ext_tnl_
id      Ingress  Egress
5001      2202      R1-R3-
Primary    Ingress   1.1.1.1          1.1.1.1  3.3.3.3
Total LSP protected : 1
Bandwidth in use : 0

#show rsvp session detail
Ingress (Primary)
3.3.3.3
  From: 1.1.1.1, LSPstate: Up, LSPname: R1-R3-Primary
  Ingress FSM state: Operational
  Establishment Time: 0s 8ms
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  IGP-Shortcut: Disabled, LSP metric: 2
  LSP Protection: facility
  Fast-Reroute bandwidth : 0
  Protection type desired: Link
  Fast-Reroute Setup priority: 7, Hold priority: 0
  Bypass trunk: B1-B4, Merge Point: 10.2.3.32, MP Label: 3
  Bypass OutLabel: 24321, OutIntf: xe5
  Protection provided -> Type: Link,BW: Best-effort
  Label in: -, Label out: 24321
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2202, Ext-Tunnel Id: 1.1.1.1
  Downstream: 10.1.2.21, xe1
  Path refresh: 30 seconds (RR enabled) (due in 26564 seconds)
  Resv refresh: 0 seconds (due in 0 seconds)
  Resv lifetime: 157 seconds (due in 150 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Enabled
  Admin Groups: none
  Configured Path: primary_1 (in use)
  Configured Explicit Route Detail :
    10.1.2.21/32 strict
    10.2.3.32/32 strict
  Session Explicit Route Detail :
    10.1.2.21/32 strict

```

```

10.2.3.32/32 strict
Record route:
LP = 1 -> PLR's Downstream link is protected      PU = 1 -> Protection is in use on PLR
NP = 1 -> PLR's Downstream neighbor is protected  BP = 1 -> BW protection available at PLR
-----
IP Address      Label      (LP, PU, NP, BP)
-----
<self>
10.1.2.21      24321      ( 0,  0,  0,  0)
10.2.3.32      3          ( 0,  0,  0,  0)
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Ingress (Bypass)
3.3.3.3
From: 1.1.1.1, LSPstate: Up, LSPname: B1-B4-Bypass
Ingress FSM state: Operational
Establishment Time: 0s 14ms
Setup priority: 7, Hold priority: 0
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
IGP-Shortcut: Disabled, LSP metric: 2
LSP Protection: None
Bypass trunk bandwidth type: Best-effort
Label in: -, Label out: 24321
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5002, LSP Id: 2203, Ext-Tunnel Id: 1.1.1.1
Downstream: 10.1.4.41, xe5
Path refresh: 30 seconds (RR enabled) (due in 25747 seconds)
Resv lifetime: 157 seconds (due in 139 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: bypass_1 (in use)
Configured Explicit Route Detail :
 10.1.4.41/32 strict
 10.3.4.34/32 strict
Session Explicit Route Detail :
 10.1.4.41/32 strict
 10.3.4.34/32 strict
Record route:
-----
IP Address      Label
-----
<self>
10.1.4.41
10.3.4.34
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Total LSP protected : 1, Bandwidth in use : 0

```

P1

```

#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes

```



```
Transit RSVP:
To          From          Type LSPName          State
Uptime     Rt  Style Labelin Labelout DSType
3.3.3.3    1.1.1.1 PRI R1-R3-Primary    UP    00:57:44  1 1
SE        24321 3      ELSP_CON
Total 1 displayed, Up 1, Down 0.
```

```
#show rsvp session de
Transit
3.3.3.3
  From: 1.1.1.1, LSPstate: Up, LSPname: R1-R3-Primary
  Transit upstream state: Operational, downstream state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: facility
  Fast-Reroute bandwidth : 0
  Protection type desired: Link
  Fast-Reroute Setup priority: 7, Hold priority: 0
  Label in: 24321, Label out: 3
  Tspec rate: 0, Espec rate: 0
  Tunnel Id: 5001, LSP Id: 2202, Ext-Tunnel Id: 1.1.1.1
  Downstream: 10.2.3.32, xe10 Upstream: 10.1.2.12, xe1
  Path refresh: 30 seconds (RR enabled) (due in 26500 seconds)
  Path lifetime: 157 seconds (due in 126 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 20926 seconds)
  Resv lifetime: 157 seconds (due in 151 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Enabled
  Admin Groups: Received Explicit Route Detail :
    10.1.2.21/32 strict
    10.2.3.32/32 strict
  Session Explicit Route Detail :
    10.2.3.32/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  10.1.2.12      24321
  <self>
  10.2.3.32      3
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type: ELSP_CONFIG
  CLASS DSCP_value EXP_value
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
```

P2

```
#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
```

```
Egress RSVP:
To          From          Type LSPName          State
Uptime     Rt  Style Labelin Labelout DSType
3.3.3.3    1.1.1.1 PRI R1-R3-Primary    UP    00:58:47  1 1
SE        3      -      ELSP_CON
3.3.3.3    1.1.1.1 PRI B1-B4-Bypass     UP    01:12:30  1 1
SE        3      -      ELSP_CON
Total 2 displayed, Up 2, Down 0
```

```
#show rsvp session detail
Egress
3.3.3.3
  From: 1.1.1.1, LSPstate: Up, LSPname: R1-R3-Primary
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: facility
  Fast-Reroute bandwidth : 0
  Protection type desired: Link
  Fast-Reroute Setup priority: 7, Hold priority: 0
  Label in:      3, Label out: -
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5001, LSP Id: 2202, Ext-Tunnel Id: 1.1.1.1
  Upstream: 10.2.3.23, xe10
  Path lifetime: 157 seconds (due in 140 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 37780 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Enabled
  Admin Groups:   Received Explicit Route Detail :
    10.2.3.32/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  10.1.2.12       24321
  10.2.3.23       3
  <self>
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type:  ELSP_CONFIG
  CLASS      DSCP_value      EXP_value
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Egress
3.3.3.3
  From: 1.1.1.1, LSPstate: Up, LSPname: B1-B4-Bypass
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in:      3, Label out: -
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5002, LSP Id: 2203, Ext-Tunnel Id: 1.1.1.1
  Upstream: 10.3.4.43, xe5
  Path lifetime: 157 seconds (due in 134 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 29222 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups:   Received Explicit Route Detail :
    10.3.4.34/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  10.1.4.14
  10.3.4.43
  <self>
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type:  ELSP_CONFIG
  CLASS      DSCP_value      EXP_value
  Last Recorded Error Code: None
```

```
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls.
```

PE2

```
#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
```

```
Transit RSVP:
To          From          Type LSPName          State
Uptime     Rt  Style  Labelin Labelout DStype
3.3.3.3      1.1.1.1      PRI   B1-B4-Bypass     UP    01:14:12  1 1
SE      24321    3      ELSP_CON
Total 1 displayed, Up 1, Down 0.
```

```
#show rsvp session detail
Transit
3.3.3.3
  From: 1.1.1.1, LSPstate: Up, LSPname: B1-B4-Bypass
  Transit upstream state: Operational, downstream state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: 24321, Label out: 3
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5002, LSP Id: 2203, Ext-Tunnel Id: 1.1.1.1
  Downstream: 10.3.4.34, xe5 Upstream: 10.1.4.14, xe2
  Path refresh: 30 seconds (RR enabled) (due in 25543 seconds)
  Path lifetime: 157 seconds (due in 146 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 17729 seconds)
  Resv lifetime: 157 seconds (due in 135 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: Received Explicit Route Detail :
    10.1.4.41/32 strict
    10.3.4.34/32 strict
  Session Explicit Route Detail :
    10.3.4.34/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  10.1.4.14
  <self>
  10.3.4.34
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type: ELSP_CONFIG
  CLASS DSCP_value EXP_value
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
```

Limitations

Dedicated Bypass Bandwidth

Refer the topology defined above.

Suppose we have two primary tunnels P1 (100mbps) and P2(20mbps) ingressing from R1 and egressing at R3 (path R1->R2>R3) and asking for BW protection and we have two Bypass tunnels bp1 (120mbps) and bp2(80mbps) type dedicated with same ingress and egress router taking Path R1->R4->R3. Below are the two cases defined in which we can observe different kinds of behavior.

1. Let the primary P1 and P2 come up.

CASE 1:

i) If the bypass bp1 (120mbps) comes up first it will give protection to both the primaries P1 and P2. bp2 should remain idle and will not give protection if there are no other primary tunnels asking for it.

CASE 2:

i) If the bypass bp2 (80mbps) comes up first it will give protection to only the primary P2 (20mbps) that will have satisfied protection which will not be changed until the bypass will go down.

ii) After that if bp1 (120mbps) comes it will provide protection to primary P1 (100mbps).

So in the CASE 1 after the protection has been provided to both the primary tunnels P1 and P2 by bypass bp1 if new primary tunnel P3 comes up with BW protection of 80mbps it would be given by bp2 (80mbps).

But in the CASE 2 as bp2 has only 60mbps left (20mbps is being used by P2) and it would not give protection to P3 tunnel and it will remain unprotected. To get the protection new tunnel has to have setup and hold priorities higher than other tunnels which are already been served with the bypass protection.

Secondary Tunnel

Suppose we have primary tunnel P1 (100mbps) ingressing from R1 and egressing at R3 (path R1->R2>R3) and asking for BW protection and we have Bypass tunnel bp1 (120mbps) type dedicated with same ingress and egress router taking Path R1->R4->R3. Then Bypass will start providing protection to primary P1.

If the primary went down it will start using the local protection.

After that if the secondary tunnel is provisioned, primary LSP, which is in using backup state shall continue to use backup path and will not shift over to secondary path.

Facility Bypass with Ring Topology Configuration

This section contains a complete Facility Bypass with Ring Topology configuration.

During facility bypass integration to OcNOS SP, few issues were reported when upstream and downstream interfaces of a session happens to be same (i.e. protection path is same as upstream path) and also CSPF most likely had some issues where LSP path used to formed by crossing the head node of the path.

Considering the information available in RSVP to impose restriction, bypass tunnel path crossing primary LSP node anywhere in between merge point were not considered for mapping.

Below assumption point was added in ERD and documents were updated on the line.

If protection is requested by primary session, then initial bypass matching criteria will be to ensure egress (merge point) node of bypass will be one of the nodes of primary LSP and bypass never intersect any node of primary LSP until the merge point.

The facility bypass method takes advantage of the MPLS label stack. Instead of creating a separate LSP for every backed-up LSP, a single LSP is created that serves to back up a set of LSPs. We call such an LSP tunnel a bypass tunnel. The bypass tunnel must intersect the path of the original LSP(s) somewhere downstream of the PLR.


Naturally, this constrains the set of LSPs being backed up via that bypass tunnel to those that pass through some common downstream node. All LSPs that pass through the point of local repair and through this common node that do not also use the facilities involved in the bypass tunnel are candidates for this set of LSPs.

By multiple facility bypass tunnels, we mean that multiple facility bypass tunnels can be created to the same egress/MP. For a protected LSP there could be multiple candidates available. The mapping of the LSP to one of the backup tunnels has to be efficiently done so that we can extract the maximum benefit out of those backup tunnels available

Topology

[Figure 28](#) displays a sample Facility Bypass with Ring topology.

Figure 28. Facility Bypass with Ring Topology



Configuration

All configuration commands in the table below should be followed for each router.

R1

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#exit | Exit interface mode. |
| (config)#bfd interval 3 minrx 3 multiplier | Configure BFD interval |
| (config)#router-id 1.1.1.1 | Assigning router-id |
| (config)#router rsvp | Enter router mode for RSVP. |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface xe8 | Specify the Interface to be configured. |
| (config-if)#ip address 10.1.1.1/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to |

| | |
|---|--|
| | point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe20 | Specify the Interface to be configured |
| (config-if)#ip address 12.1.1.1/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis ISIS-IGP | Create an IS-IS routing instance |
| (config-router)#is-type level-1 | Configure instance as level-1 only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng router-id 1.1.1.1 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1 |
| (config-router)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config-router)#bfd all-interfaces | Enable BFD for all neighbors. |
| (config-router)#net 49.0000.0000.0001.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)#exit | Exit router mode. |
| (config)#rsvp-path R1-R5-PRI-001 | Create a rsvp path |
| (config-path)#10.1.1.2 strict | Configure this explicit router path as a strict hop |
| (config-path)#14.1.1.3 strict | Configure this explicit router path as a strict hop |
| (config-path)#17.1.1.3 strict | Configure this explicit router path as a strict hop |
| (config-path)#exit | Exit the rsvp-path mode |
| (config)#rsvp-path R1-R5-BPS-001 | Create a rsvp path |
| (config-path)#12.1.1.2 strict | Configure this explicit router path as a strict hop |
| (config-path)#15.1.1.3 strict | Configure this explicit router path as a strict hop |
| (config-path)#exit | Exit the rsvp-path mode |
| (config)#rsvp-trunk R1-R5-PRI-001 | Enter the trunk mode for RSVP |
| (config-trunk)#primary fast-reroute protection facility | Configure primary fast-reroute protection facility for a trunk. |

| | |
|---|---|
| (config-trunk)#primary fast-reroute node-protection | Configure primary fast-reroute node protection for a trunk. |
| (config-trunk)#primary path R1-R5-PRI-001 | Configure trunk to use the defined path. |
| (config-trunk)#to 5.5.5.5 | Specify the IPv4 egress (destination point) for the LSP |
| (config-path)#exit | Exit the rsvp-trunk mode |
| (config)#rsvp-bypass R1-R5-BPS-001 | Enter the bypass mode for RSVP |
| (config-trunk)#path R1-R5-BPS-001 | Configure path for bypass tunnel |
| (config-trunk)#to 5.5.5.5 | Specify the IPv4 egress (destination point) for the LSP |
| (config-path)#exit | Exit the rsvp-bypass mode |
| (config)#commit | Commit the transaction. |

R2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 2.2.2.2 | Assigning router-id |
| (config)#bfd interval 3 minrx 3 multiplier | Configure BFD interval |
| (config)#router rsvp | Enter router mode for RSVP. |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface xe29 | Specify the Interface to be configured. |
| (config-if)#ip address 10.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe5 | Specify the Interface to be configured |
| (config-if)#ip address 14.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |

| | |
|--|--|
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis ISIS-IGP | Create an IS-IS routing instance |
| (config-router)#is-type level-1 | Configure instance as level-1 only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng router-id 2.2.2.2 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1 |
| (config-router)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config-router)#bfd all-interfaces | Enable BFD for all neighbors. |
| (config-router)#net 49.0000.0000.0002.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)#exit | Exit router mode. |
| (config)#commit | Commit the transaction. |

R3

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 3.3.3.3 | Assigning router-id |
| (config)#bfd interval 3 minrx 3 multiplier | Configure BFD interval |
| (config)#router rsvp | Enter router mode for RSVP. |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface xe7 | Specify the Interface to be configured. |
| (config-if)#ip address 12.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe6 | Specify the Interface to be configured |

| | |
|--|--|
| (config-if)#ip address 15.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis ISIS-IGP | Create an IS-IS routing instance |
| (config-router)#is-type level-1 | Configure instance as level-1 only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng router-id 3.3.3.3 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1 |
| (config-router)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config-router)#bfd all-interfaces | Enable BFD for all neighbors. |
| (config-router)#net 49.0000.0000.0003.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)#exit | Exit router mode. |
| (config)#commit | Commit the transaction. |

R4

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 4.4.4.4 | Assigning router-id |
| (config)#bfd interval 3 minrx 3 multiplier | Configure BFD interval |
| (config)#router rsvp | Enter router mode for RSVP. |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface xe4 | Specify the Interface to be configured. |
| (config-if)#ip address 17.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |

| | |
|--|--|
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe6 | Specify the Interface to be configured |
| (config-if)#ip address 14.1.1.3/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis ISIS-IGP | Create an IS-IS routing instance |
| (config-router)#is-type level-1 | Configure instance as level-1 only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng router-id 4.4.4.4 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1 |
| (config-router)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config-router)#bfd all-interfaces | Enable BFD for all neighbors. |
| (config-router)#net 49.0000.0000.0004.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)#exit | Exit router mode. |
| (config)#commit | Commit the transaction. |

R5

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 5.5.5.5/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 5.5.5.5 | Assigning router-id |
| (config)#bfd interval 3 minrx 3 multiplier | Configure BFD interval |
| (config)#router rsvp | Enter router mode for RSVP. |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface xe12 | Specify the Interface to be configured. |

| | |
|--|--|
| (config-if)#ip address 17.1.1.3/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xel | Specify the Interface to be configured |
| (config-if)#ip address 15.1.1.3/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-rsvp | Enable rsvp configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis ISIS-IGP | Create an IS-IS routing instance |
| (config-router)#is-type level-1 | Configure instance as level-1 only routing. |
| (config-router)#metric-style wide | Configure the new style of metric type as wide. |
| (config-router)#mpls traffic-eng router-id 5.5.5.5 | Configure MPLS-TE unique router-id TLV. |
| (config-router)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1 |
| (config-router)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config-router)#bfd all-interfaces | Enable BFD for all neighbors. |
| (config-router)#net 49.0000.0000.0005.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)#exit | Exit router mode. |
| (config)#commit | Commit the transaction. |

Validation

RSVP Session

Validate that the RSVP Session is up.

R1:

```
R1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
```

State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
 * indicates the session is active with local repair at one or more nodes
 (P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Type | LSPName | State |
|-----------|---------------|----------|-----------------------|----------------------|
| Uptime Rt | Style Labelin | Labelout | DSType | |
| 5.5.5.5 | 1.1.1.1 | PRI | R1-R5-PRI-001-Primary | UP 00:49:18 1 1 SE - |
| | 52480 | DEFAULT | | |
| 5.5.5.5 | 1.1.1.1 | BPS | R1-R5-BPS-001-Bypass | UP 05:24:23 1 1 SE - |
| | 25600 | DEFAULT | | |

Total 2 displayed, Up 2, Down 0.

R1#show rsvp session detail

Ingress (Primary)

5.5.5.5

From: 1.1.1.1, LSPstate: Up, LSPname: R1-R5-PRI-001-Primary
 Ingress FSM state: Operational
 Establishment Time: 322s 925ms
 Setup priority: 7, Hold priority: 0
 CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
 LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: ISIS
 IGP-Shortcut: Disabled, LSP metric: 30
 LSP Protection: facility
 Fast-Reroute bandwidth : 0
 Protection type desired: Node
 Fast-Reroute Hop limit: 255
 Fast-Reroute Setup priority: 7, Hold priority: 0
 Bypass trunk: R1-R5-BPS-001, Merge Point: 17.1.1.3, MP Label: 25600
 Bypass OutLabel: 25600, OutIntf: xe20
 Protection provided -> Type: Node, BW: Best-effort
 Label in: -, Label out: 52480,
 Tspec rate: 0, Fspec rate: 0
 Policer: Not Configured
 Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 1.1.1.1
 Bind value: 0, Oper state: NA, Alloc mode: NA
 Downstream: 10.1.1.2, xe8
 Path refresh: 30 seconds (RR enabled) (due in 27023 seconds)
 Resv refresh: 0 seconds (due in 0 seconds)
 Resv lifetime: 157 seconds (due in 128 seconds)
 Retry count: 0, intrvl: 30 seconds
 RRO re-use as ERO: Disabled
 Label Recording: Enabled
 Admin Groups: none
 Configured Path: R1-R5-PRI-001 (in use)
 Configured Explicit Route Detail :
 10.1.1.2/32 strict
 14.1.1.3/32 strict
 17.1.1.3/32 strict
 Session Explicit Route Detail :
 10.1.1.2/32 strict
 14.1.1.3/32 strict
 17.1.1.3/32 strict
 Record route:
 LP = 1 -> PLR's Downstream link is protected PU = 1 -> Protection is in use on PLR
 NP = 1 -> PLR's Downstream neighbor is protected BP = 1 -> BW protection available at PLR

| IP Address | Label | (LP, PU, NP, BP) |
|------------|-------|------------------|
| <self> | | |
| 10.1.1.2 | 52480 | (0, 0, 0, 0) |
| 14.1.1.3 | 52480 | (0, 0, 0, 0) |
| 17.1.1.3 | 25600 | (0, 0, 0, 0) |

 Style: Shared Explicit Filter
 Traffic type: controlled-load
 Minimum Path MTU: 9216
 Recorded Time : N/A
 Current Error:

```

    Code : None, Value : None
    Originated Node : None, Recorded Time : N/A
    Last Signaled Error:
    Code : None, Value : None
    Originated Node : None, Recorded Time : N/A
    Trunk Type: mpls
    Ingress (Bypass)
5.5.5.5
    From: 1.1.1.1, LSPstate: Up, LSPname: R1-R5-BPS-001-Bypass
    Ingress FSM state: Operational
    Establishment Time: 0s 4ms
    Setup priority: 7, Hold priority: 0
    CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
    LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: ISIS
    IGP-Shortcut: Disabled, LSP metric: 20
    LSP Protection: None
    Bypass trunk bandwidth type: Best-effort
    Label in: -, Label out: 25600,
    Tspec rate: 0, Fspec rate: 0
    Policer: Not Configured
    Tunnel Id: 5002, LSP Id: 2205, Ext-Tunnel Id: 1.1.1.1
    Bind value: 0, Oper state: NA, Alloc mode: NA
    Downstream: 12.1.1.2, xe20
    Path refresh: 30 seconds (RR enabled) (due in 10514 seconds)
    Resv lifetime: 157 seconds (due in 141 seconds)
    Retry count: 0, intrvl: 30 seconds
    RRO re-use as ERO: Disabled
    Label Recording: Disabled
    Admin Groups: none
    Configured Path: R1-R5-BPS-001 (in use)
    Configured Explicit Route Detail :
    12.1.1.2/32 strict
    15.1.1.3/32 strict
    Session Explicit Route Detail :
    12.1.1.2/32 strict
    15.1.1.3/32 strict
    Record route:
    -----
    IP Address      Label
    -----
    <self>
    12.1.1.2
    15.1.1.3
    Style: Shared Explicit Filter
    Traffic type: controlled-load
    Minimum Path MTU: 9216
    Recorded Time : N/A
    Current Error:
    Code : None, Value : None
    Originated Node : None, Recorded Time : N/A
    Last Signaled Error:
    Code : RSVP System error (23), Value : N/A (0)
    Originated Node : 15.1.1.3, Recorded Time : 2023 May 16 08:52:51
    Trunk Type: mpls
    Total LSP protected : 1, Bandwidth in use : 0

```

R2:

```

R2#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Transit RSVP:
To          From          Type    LSPName          State
Uptime    Rt  Style  Labelin  Labelout  DSType

```

```

5.5.5.5      1.1.1.1      PRI      R1-R5-PRI-001-Primary      UP      00:49:59  1 1
SE      52480      52480      ELSP_CON
Total 1 displayed, Up 1, Down 0.

```

R2#show rsvp session detail

Transit

5.5.5.5

```

From: 1.1.1.1, LSPstate: Up, LSPname: R1-R5-PRI-001-Primary
Transit upstream state: Operational, downstream state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: facility
Fast-Reroute bandwidth : 0
Protection type desired: Node
Fast-Reroute Hop limit: 255
Fast-Reroute Setup priority: 7, Hold priority: 0
Label in: 52480, Label out: 52480,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 1.1.1.1
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 14.1.1.3, xe5 Upstream: 10.1.1.1, xe29
Path refresh: 30 seconds (RR enabled) (due in 27004 seconds)
Path lifetime: 157 seconds (due in 130 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 19943 seconds)
Resv lifetime: 157 seconds (due in 141 seconds)
RRO re-use as ERO: Disabled
Label Recording: Enabled
Admin Groups: Received Explicit Route Detail :
  10.1.1.2/32 strict
  14.1.1.3/32 strict
  17.1.1.3/32 strict
Session Explicit Route Detail :
  14.1.1.3/32 strict
  17.1.1.3/32 strict
Record route:
-----
IP Address      Label
-----
10.1.1.1        52480
<self>
14.1.1.3        52480
17.1.1.3        25600
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 9216
LSP Type: ELSP_CONFIG
CLASS   DSCP_value   EXP_value
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls

```

R3:

R3#show rsvp session

```

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

Transit RSVP:

```

To      From      Type      LSPName      State
Uptime  Rt  Style  Labelin  Labelout  DSType
5.5.5.5  1.1.1.1  PRI      R1-R5-BPS-001-Bypass  UP      05:25:48  1 1
SE      25600  3      ELSP_CON
Total 1 displayed, Up 1, Down 0.

```

```

R3#show rsvp session detail
Transit
5.5.5.5
  From: 1.1.1.1, LSPstate: Up, LSPname: R1-R5-BPS-001-Bypass
  Transit upstream state: Operational, downstream state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: 25600, Label out: 3,
  Tspeg rate: 0, Fspeg rate: 0
  Tunnel Id: 5002, LSP Id: 2205, Ext-Tunnel Id: 1.1.1.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 15.1.1.3, xe6 Upstream: 12.1.1.1, xe7
  Path refresh: 30 seconds (RR enabled) (due in 10445 seconds)
  Path lifetime: 157 seconds (due in 155 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 24008 seconds)
  Resv lifetime: 157 seconds (due in 140 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: Received Explicit Route Detail :
    12.1.1.2/32 strict
    15.1.1.3/32 strict
  Session Explicit Route Detail :
    15.1.1.3/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  12.1.1.1
  <self>
  15.1.1.3
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 9216
  LSP Type: ELSP_CONFIG
  CLASS DSCP_value EXP_value
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
R3#

```

R4:

```

R4#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Transit RSVP:
To      From      Type      LSPName      State
Uptime  Rt  Style  Labelin  Labelout  DSType
5.5.5.5  1.1.1.1  PRI      R1-R5-PRI-001-Primary  UP    00:51:13  1 1
SE      52480  25600  ELSP_CON
Total 1 displayed, Up 1, Down 0.

R4#show rsvp session detail
Transit
5.5.5.5
  From: 1.1.1.1, LSPstate: Up, LSPname: R1-R5-PRI-001-Primary
  Transit upstream state: Operational, downstream state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65

```

```

LSP Protection: facility
Fast-Reroute bandwidth : 0
Protection type desired: Node
Fast-Reroute Hop limit: 255
Fast-Reroute Setup priority: 7, Hold priority: 0
Label in: 52480, Label out: 25600,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 1.1.1.1
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 17.1.1.3, xe4 Upstream: 14.1.1.2, xe6
Path refresh: 30 seconds (RR enabled) (due in 26908 seconds)
Path lifetime: 157 seconds (due in 148 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 37164 seconds)
Resv lifetime: 157 seconds (due in 144 seconds)
RRO re-use as ERO: Disabled
Label Recording: Enabled
Admin Groups: Received Explicit Route Detail :
  14.1.1.3/32 strict
  17.1.1.3/32 strict
Session Explicit Route Detail :
  17.1.1.3/32 strict
Record route:
-----
IP Address      Label
-----
10.1.1.1        52480
14.1.1.2        52480
<self>
17.1.1.3        25600
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 9216
LSP Type: ELSP_CONFIG
CLASS    DSCP_value    EXP_value
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls

```

R5:

```

R5#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

```

Egress RSVP:
To      From      Type      LSPName      State
Uptime  Rt  Style  Labelin  Labelout  DStype
5.5.5.5      1.1.1.1      PRI      R1-R5-PRI-001-Primary  UP    00:51:45  1 1
SE      25600      -      ELSP_CON
5.5.5.5      1.1.1.1      PRI      R1-R5-BPS-001-Bypass  UP    05:26:50  1 1
SE      3      -      ELSP_CON
Total 2 displayed, Up 2, Down 0.

```

```

R5#show rsvp session detail
Egress
5.5.5.5
  From: 1.1.1.1, LSPstate: Up, LSPname: R1-R5-PRI-001-Primary
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: facility
  Fast-Reroute bandwidth : 0
  Protection type desired: Node

```



```

Fast-Reroute Hop limit: 255
Fast-Reroute Setup priority: 7, Hold priority: 0
Label in: 25600, Label out: -,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 1.1.1.1
Bind value: 0, Oper state: NA, Alloc mode: NA
Upstream: 17.1.1.2, xe12
Path lifetime: 157 seconds (due in 126 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 28434 seconds)
RRO re-use as ERO: Disabled
Label Recording: Enabled
Admin Groups: Received Explicit Route Detail :
  17.1.1.3/32 strict
Record route:
-----
IP Address      Label
-----
10.1.1.1        52480
14.1.1.2        52480
17.1.1.2        25600
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 9216
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Egress
5.5.5.5
From: 1.1.1.1, LSPstate: Up, LSPname: R1-R5-BPS-001-Bypass
Egress FSM state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 3, Label out: -,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5002, LSP Id: 2205, Ext-Tunnel Id: 1.1.1.1
Bind value: 0, Oper state: NA, Alloc mode: NA
Upstream: 15.1.1.2, xe1
Path lifetime: 157 seconds (due in 141 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 927 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: Received Explicit Route Detail :
  15.1.1.3/32 strict
Record route:
-----
IP Address      Label
-----
12.1.1.1
15.1.1.2
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 9216
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls

```

RSVP Bypass

Validate that the RSVP bypass session is up.

R1:

```
R1#show rsvp bypass
Ingress RSVP:
To          From          LSPName          State
Uptime    Rt   Style  Labelin  Labelout  DStype
5.5.5.5    25600 1.1.1.1 R1-R5-BPS-001-Bypass UP      05:27:41 1 1 SE  -
                DEFAULT
```

To validate RSVP bypass session details:

```
R1#show rsvp bypass detail
Ingress (Bypass)
5.5.5.5
  From: 1.1.1.1, LSPstate: Up, LSPname: R1-R5-BPS-001-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 4ms
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out: 25600,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5002, LSP Id: 2205, Ext-Tunnel Id: 1.1.1.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 12.1.1.2, xe20
  Path refresh: 30 seconds (RR enabled) (due in 10319 seconds)
  Resv lifetime: 157 seconds (due in 126 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: R1-R5-BPS-001 (in use)
  Configured Explicit Route Detail :
    12.1.1.2/32 strict
    15.1.1.3/32 strict
  Session Explicit Route Detail :
    12.1.1.2/32 strict
    15.1.1.3/32 strict
  Record route:
  -----
  IP Address          Label
  -----
  <self>
  12.1.1.2
  15.1.1.3
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 9216
  Recorded Time : N/A
  Current Error:
    Code : None, Value : None
    Originated Node : None, Recorded Time : N/A
  Last Signaled Error:
    Code : RSVP System error (23), Value : N/A (0)
    Originated Node : 15.1.1.3, Recorded Time : 2023 May 16 08:52:51
  Trunk Type: mpls
  Total LSP protected : 1, Bandwidth in use : 0
```

To validate RSVP bypass Protected-lsp-list:


```
R1#show rsvp bypass protected-lsp-list
Bypass trunk: R1-R5-BPS-001
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id      Lsp Id      Lsp Name      Role      Ext_tnl_
id      Ingress      Egress
5001      2201      R1-R5-PRI-001-
Primary      Ingress      1.1.1.1      1.1.1.1      5.5.5.5
Total LSP protected : 1
Bandwidth in use : 0
```

RSVP-Multipath Configuration

RSVP multipath provides multiple paths to reach the destination. RSVP services (L2 VPN, L3 VPN, LU, mapped route, etc.) can be mapped to RSVP multipath so that traffic is load-balanced. At ingress node the traffic is load-balanced based on the configured hash (L3 src/dest IP/port, L2 src/dst mac, or in-label if intermediate Autonomous segments). Each LSP path within multipath group can provide individual protection for each path (facility, 1-to-1, secondary). Each tunnel path in multipath group cost may vary (can have different hop, with consideration for load-balancing the traffic). Each member can have the backup protection as other multipath member path or outside the multipath member path.

Topology

Figure 29. RSVP-Multipath Topology



Configuration

R1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#hardware-profile statistics tunnel-lif enable | Configure hardware-profile statistics |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 40.40.40.40/32 secondary | Set a secondary IP address of the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#router-id 40.40.40.40 | Set the router ID to IP address 40.40.40.40 |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ip address 10.10.1.1/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |

| | |
|---|---|
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xel1 | Enter interface mode. |
| (config-if)#ip address 13.1.1.1/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the Routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 40.40.40.40 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 10.1.1.0/24 area 0 | Define the network (10.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 13.1.1.0/24 area 0 | Define the network (13.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 40.40.40.40/32 area 0 | Set the IP address of the loopback interface to 40.40.40.40/32. |
| (config-router)#exit | Exit Router mode. |
| (config)#rsvp-path path1 mpls | Enter the path mode for RSVP path1. |
| (config-path)#10.1.1.2 strict | Configure this explicit route path as a strict hop. |
| (config-path)#11.1.1.3 strict | Configure this explicit route path as a strict hop. |
| (config-path)#12.1.1.4 strict | Configure this explicit route path as a strict hop. |
| (config)#exit | Exit the path mode. |
| (config)#rsvp-path path2 mpls | Enter the path mode for RSVP path2. |
| (config-path)#13.1.1.2 strict | Configure this explicit route path as a strict hop. |
| (config-path)#14.1.1.3 strict | Configure this explicit route path as a strict hop. |
| (config-path)#15.1.1.4 strict | Configure this explicit route path as a strict hop. |
| (config)#exit | Exit the path mode. |
| (config)#rsvp-trunk t1 ipv4 | Enter the trunk mode for rsvp. |
| (config-trunk)#primary path path1 | Configure trunk t1 to use the defined path. |
| (config-trunk)#from 40.40.40.40 | Specify the IPv4 ingress (source point) for the LSP. |
| (config-trunk)#to 47.47.47.47 | Specify the IPv4 egress (destination point) for the LSP. |
| (config-trunk)#exit | Exit from trunk mode. |
| (config)#rsvp-trunk t2 ipv4 | Enter the trunk mode for rsvp. |
| (config-trunk)#primary path path2 | Configure trunk t2 to use the defined path. |
| (config-trunk)#from 40.40.40.40 | Specify the IPv4 egress (source point) for the LSP. |

| | |
|---|---|
| (config-trunk)#to 47.47.47.47 | Specify the IPv4 ingress (destination point) for the LSP. |
| (config-trunk)#exit | Exit from trunk mode. |
| (config)#rsvp-multipath test | Configure RSVP Multipath group. |
| (config-multipath)#to 47.47.47.47 | Configure a mandatory destination FEC(egress router). |
| (config-multipath)#description "my-group" | Configure description to RSVP Multipath group. |
| (config-multipath)#exit | Exit from RSVP Multipath mode. |
| (config)#rsvp-trunk t1 ipv4 | Enter the trunk mode for rsvp. |
| (config-trunk)#multipath-group test | Configure RSVP Multipath group to trunk t1. |
| (config-trunk)#exit | Exit from trunk mode. |
| (config)#rsvp-trunk t2 ipv4 | Enter the trunk mode for rsvp. |
| (config-trunk)#multipath-group test | Configure RSVP Multipath group to trunk t2. |
| (config-trunk)#exit | Exit from trunk mode. |
| (config)#commit | Commit all the configurations. |

R2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 46.46.46.46/32 secondary | Set a secondary IP address of the interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xel | Enter interface mode. |
| (config-if)#ip address 10.10.1.2/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe5 | Enter interface mode. |
| (config-if)#ip address 11.1.1.2/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the Routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 46.46.46.46 | Configure OSPF router-ID same as loopback interface IP address |

| | |
|---|--|
| (config-router)#network 10.1.1.0/24 area 0 | Define the network (10.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 11.1.1.0/24 area 0 | Define the network (11.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 46.46.46.46/32 area 0 | Set the IP address of the loopback interface to 46.46.46.46/32. |
| (config-router)#exit | Exit Router mode. |
| (config)#commit | Commit all the configurations. |

R3

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 48.48.48.48/32 secondary | Set a secondary IP address of the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ip address 12.1.1.3/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe5 | Enter interface mode. |
| (config-if)#ip address 11.1.1.3/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the Routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 48.48.48.48 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 11.1.1.0/24 area 0 | Define the network (11.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 12.1.1.0/24 area 0 | Define the network (12.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 48.48.48.48/32 area 0 | Set the IP address of the loopback interface to 48.48.48.48/32. |
| (config-router)#exit | Exit Router mode. |
| (config)#commit | Commit all the configurations. |

R4

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 56.56.56.56/32 secondary | Set a secondary IP address of the interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#ip address 13.1.1.2/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe17 | Enter interface mode. |
| (config-if)#ip address 14.1.1.2/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the Routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 56.56.56.56 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 13.1.1.0/24 area 0 | Define the network (13.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 14.1.1.0/24 area 0 | Define the network (14.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 56.56.56.56/32 area 0 | Set the IP address of the loopback interface to 56.56.56.56/32. |
| (config-router)#exit | Exit Router mode. |
| (config)#commit | Commit all the configurations. |

R5

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 58.58.58.58/32 secondary | Set a secondary IP address of the interface |
| (config-if)#exit | Exit interface mode. |

| | |
|---|---|
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#ip address 15.1.1.3/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe17 | Enter interface mode. |
| (config-if)#ip address 14.1.1.3/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the Routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 58.58.58.58 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 15.1.1.0/24 area 0 | Define the network (15.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 14.1.1.0/24 area 0 | Define the network (14.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 58.58.58.58/32 area 0 | Set the IP address of the loopback interface to 58.58.58.58/32. |
| (config-router)#exit | Exit Router mode. |
| (config)#commit | Commit all the configurations. |

R6

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 47.47.47.47/32 secondary | Set a secondary IP address of the interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ip address 12.1.1.4/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |

| | |
|---|---|
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#ip address 15.1.1.4/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the Routing process and specify the Process ID (1). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 47.47.47.47 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 12.1.1.0/24 area 0 | Define the network (12.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 15.1.1.0/24 area 0 | Define the network (15.1.1.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 47.47.47.47/32 area 0 | Set the IP address of the loopback interface to 47.47.47.47/32. |
| (config-router)#exit | Exit Router mode. |
| (config)#commit | Commit all the configurations. |

Validation

R1

```
#show ip ospf neighbor
```

```
Total number of full neighbors: 2
```

```
OSPF process 1 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|----------|-----------|-------------|
| 46.46.46.46 | 1 | Full/DR | 00:00:38 | 10.1.1.2 | ce54 | 0 |
| 56.56.56.56 | 1 | Full/DR | 00:00:37 | 13.1.1.2 | xe11 | 0 |

```
#show rsvp session
```

```
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
```

```
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
```

```
* indicates the session is active with local repair at one or more nodes
```

```
(P) indicates the secondary-priority session is acting as primary
```

```
Ingress RSVP:
```

| To | From | Type | LSPName | State | |
|-------------|-------------|-------|------------|----------|----------|
| Uptime | Rt | Style | Labelin | Labelout | DSType |
| 47.47.47.47 | 40.40.40.40 | PRI | t1-Primary | UP | 00:01:46 |
| 24320 | DEFAULT | | | | 1 1 SE |
| 47.47.47.47 | 40.40.40.40 | PRI | t2-Primary | UP | 00:01:46 |
| 24320 | DEFAULT | | | | 1 1 SE |

```
Total 2 displayed, Up 2, Down 0.
```

```
#show rsvp multipath
```

```
RSVP-multipath Name : test, ID : 101
```

```
Description : "my-group"
```

```
Member count : 2, Egress : 47.47.47.47/32
```

```
Member details :
```

```

-----
Trunk-ID      Trunk-name      Status
5001          t1                active
5002          t2                active

#show mpls rsvp-multipath
Codes: > - installed FTN, * - selected FTN, t - tunnel, R - RSVP-TE FTN

Multipath Name : test, ID : 101, Nhlfe Ix : 4
Active member count : 2, FEC : 47.47.47.47/32
Active member details :
-----
Index  Code      FTN-ID  Nhlfe-ID  Tunnel-id  Pri  LSP-Type  Out-Label  Out-
Intf   ELC      Nexthop
1      R(t)>    1        3          5001      Yes  LSP_
DEFAULT 24320    ce54      No         10.1.1.2
2      R(t)>    2        5          5002      Yes  LSP_
DEFAULT 24320    xe11      No         13.1.1.2

#show mpls counters rsvp multipath-name test
Tunnel-id 5001 Extended Tunnel-ID 40.40.40.40 Egress 47.47.47.47
lsp-name : t1-Primary [Ingress]
lsp-ingress : 40.40.40.40      lsp-id : 2201
Rx pkts : n/a                Rx bytes : n/a
Tx pkts : 864364             Tx bytes : 1298276230

Tunnel-id 5002 Extended Tunnel-ID 40.40.40.40 Egress 47.47.47.47
lsp-name : t2-Primary [Ingress]
lsp-ingress : 40.40.40.40      lsp-id : 2202
Rx pkts : n/a                Rx bytes : n/a
Tx pkts : 864366             Tx bytes : 1298277732

#show mpls ftn-table
Primary FTN entry with FEC: 47.47.47.47/32, id: 1, row status: Active, Tunnel-Policy: N/A
Owner: RSVP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 5001, Protected LSP id: 2201, QoS Resource id: 2, Description: t1, Color: 0
Multipath group: test
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
Owner: RSVP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: RSVP, Stale: NO, out intf: ce54, out label: 24320
Nexthop addr: 10.1.1.2      cross connect ix: 3, op code: Push

Primary FTN entry with FEC: 47.47.47.47/32, id: 2, row status: Active, Tunnel-Policy: N/A
Owner: RSVP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 5002, Protected LSP id: 2202, QoS Resource id: 3, Description: t2, Color: 0
Multipath group: test
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 5
Owner: RSVP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 5, owner: RSVP, Stale: NO, out intf: xe11, out label: 24320
Nexthop addr: 13.1.1.2      cross connect ix: 4, op code: Push

#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
      B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
      L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
      U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport

```

```

Code      FEC      FTN-ID  Nhlfe-ID  Tunnel-id  Pri  LSP-Type  Out-Label  Out-
Intf      ELC      Nexthop
R(t)>    47.47.47.47/32  1        3          5001      Yes  LSP_
DEFAULT 24320    ce54      No         10.1.1.2
R(t)>    47.47.47.47/32  2        5          5002      Yes  LSP_
DEFAULT 24320    xe11      No         13.1.1.2

```

```
#show rsvp session detail
Egress
40.40.40.40
  From: 47.47.47.47, LSPstate: Up, LSPname: PE2-PE1-Primary
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: 24960, Label out: -,
  Tspec rate: 0k, Fspec rate: 0k
  Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 47.47.47.47
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Upstream: 13.1.1.2, xell
  Path lifetime: 157 seconds (due in 146 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 34571 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: Received Explicit Route Detail :
    13.1.1.1/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  15.1.1.4
  14.1.1.3
  13.1.1.2
  <self>
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type: ELSP_CONFIG
  CLASS    DSCP_value    EXP_value
  Current Error:
    Code : None, Value : None
    Originated Node : None, Recorded Time : N/A
  Trunk Type: mpls
Ingress (Primary)
47.47.47.47
  From: 40.40.40.40, LSPstate: Up, LSPname: t1-Primary
  Ingress FSM state: Operational
  Establishment Time: 0s 22ms
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
  IGP-Shortcut: Disabled, LSP metric: 3
  LSP Protection: None
  Label in: -, Label out: 24320,
  Tspec rate: 0k, Fspec rate: 0k
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 40.40.40.40
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 10.1.1.2, ce54
  Path refresh: 30 seconds (RR enabled) (due in 29859 seconds)
  Resv lifetime: 157 seconds (due in 149 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: path1 (in use)
  Configured Explicit Route Detail :
    10.1.1.2/32 strict
    11.1.1.3/32 strict
    12.1.1.4/32 strict
  Session Explicit Route Detail :
    10.1.1.2/32 strict
    11.1.1.3/32 strict
    12.1.1.4/32 strict
  Record route:
```

```

-----
IP Address      Label
-----
<self>
10.1.1.2
11.1.1.3
12.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Ingress (Primary)
47.47.47.47
From: 40.40.40.40, LSPstate: Up, LSPname: t2-Primary
Ingress FSM state: Operational
Establishment Time: 0s 30ms
Setup priority: 7, Hold priority: 0
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
IGP-Shortcut: Disabled, LSP metric: 3
LSP Protection: None
Label in: -, Label out: 24320,
Tspec rate: 0k, Fspec rate: 0k
Policer: Not Configured
Tunnel Id: 5002, LSP Id: 2202, Ext-Tunnel Id: 40.40.40.40
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 13.1.1.2, xell
Path refresh: 30 seconds (RR enabled) (due in 29850 seconds)
Resv lifetime: 157 seconds (due in 146 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: path2 (in use)
Configured Explicit Route Detail :
  13.1.1.2/32 strict
  14.1.1.3/32 strict
  15.1.1.4/32 strict
Session Explicit Route Detail :
  13.1.1.2/32 strict
  14.1.1.3/32 strict
  15.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
<self>
13.1.1.2
14.1.1.3
15.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls

```

```
#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

Code   FEC/VRF/L2CKT   ILM-ID   In-Label   Out-Label   In-Intf   Out-
Intf/VRF   Nexthop       LSP-Type
B>   vrf-PE1         1       24320      Nolabel     N/A       vrf-
PE1      N/A         LSP_DEFAULT
R>   40.40.40.40/32   2       24960      Nolabel     N/A       N/A       127.0.0.1
      ELSP_CONFIG

R2#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID  Pri State      Dead Time  Address    Interface  Instance ID
40.40.40.40  1 Full/Backup 00:00:37  10.10.1.1  xe1        0
48.48.48.48  1 Full/Backup 00:00:34  11.1.1.3   xe12       0

R2#show rsvp session

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
Transit RSVP:
To From Type LSPName State Uptime Rt Style Labelin Labelout DStype
47.47.47.47 40.40.40.40 PRI t1-Primary UP 00:14:13 1 1 SE 25600 25600 ELSP_CON
Total 1 displayed, Up 1, Down 0.
R2#
R2#
R2#
R2#show rsvp session detail
Transit
47.47.47.47
From: 40.40.40.40, LSPstate: Up, LSPname: t1-Primary
Transit upstream state: Operational, downstream state: Operational

Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 25600, Label out: 25600,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 40.40.40.40
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 11.1.1.3, xe12 Upstream: 10.10.1.1, xe1
Path refresh: 30 seconds (RR enabled) (due in 29011 seconds)
Path lifetime: 157 seconds (due in 133 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 32105 seconds)
Resv lifetime: 157 seconds (due in 133 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: Received Explicit Route Detail :
10.10.1.2/32 strict
11.1.1.3/32 strict
12.1.1.4/32 strict
Session Explicit Route Detail :
11.1.1.3/32 strict
12.1.1.4/32 strict
Record route:
-----
```

```

IP Address Label
-----
10.10.1.1
<self>
11.1.1.3
12.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS DSCP_value EXP_value
Recorded Time : N/A
Current Error:
Code : None, Value : None
Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
R2#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
S - SNMP, L - LDP, R - RSVP, C - CRLDP
B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
P - SR Policy, U - unknown
LDP ilm-ecmp - disabled
Code FEC/VRF/L2CKT ILM-ID In-Label Out-Label In-Intf Out-Intf/VRF Nexthop pri LSP-Type
R> 47.47.47.47/32 1 25600 25600 N/A xe12 11.1.1.3 Yes ELSP_CONFIG
R3#show ip ospf neighbor
Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID Pri State Dead Time Address Interface Instance ID
46.46.46.46 1 Full/DR 00:00:31 11.1.1.2 xe40 0
47.47.47.47 1 Full/Backup 00:00:32 12.1.1.4 xe6
R3#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
Transit RSVP:
To From Type LSPName State Uptime Rt Style Labelin Labelout DStype
47.47.47.47 40.40.40.40 PRI t1-Primary UP 00:20:52 1 1 SE 25600 25600 ELSP_CON
Total 1 displayed, Up 1, Down 0.
R3#show rsvp session detail
Transit
47.47.47.47
From: 40.40.40.40, LSPstate: Up, LSPname: t1-Primary
Transit upstream state: Operational, downstream state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 25600, Label out: 25600,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 40.40.40.40
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 12.1.1.4, xe1 Upstream: 11.1.1.2, xe5
Path refresh: 30 seconds (RR enabled) (due in 28622 seconds)
Path lifetime: 157 seconds (due in 151 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 41536 seconds)
Resv lifetime: 157 seconds (due in 141 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: Received Explicit Route Detail :
11.1.1.3/32 strict
12.1.1.4/32 strict
Session Explicit Route Detail :
12.1.1.4/32 strict
Record route:
-----
IP Address Label

```

```

-----
10.10.1.1
11.1.1.2
<self>
12.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS DSCP_value EXP_value
Recorded Time : N/A
Current Error:
Code : None, Value : None
Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
R3# R3#show rsvp session detail
Transit
47.47.47.47
From: 40.40.40.40, LSPstate: Up, LSPname: t1-Primary
Transit upstream state: Operational, downstream state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 25600, Label out: 25600,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 40.40.40.40
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 12.1.1.4, xe6 Upstream: 11.1.1.2, xe40
Path refresh: 30 seconds (RR enabled) (due in 28622 seconds)
Path lifetime: 157 seconds (due in 151 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 41536 seconds)
Resv lifetime: 157 seconds (due in 141 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: Received Explicit Route Detail :
11.1.1.3/32 strict
12.1.1.4/32 strict
Session Explicit Route Detail :
12.1.1.4/32 strict
Record route:
-----
IP Address Label
-----
10.10.1.1
11.1.1.2
<self>
12.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS DSCP_value EXP_value
Recorded Time : N/A
Current Error:
Code : None, Value : None
Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
R3#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
S - SNMP, L - LDP, R - RSVP, C - CRLDP
B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
P - SR Policy, U - unknown
LDP ilm-ecmp - disabled
Code FEC/VRF/L2CKT ILM-ID In-Label Out-Label In-Intf Out-Intf/VRF Nexthop pri LSP-Type
R> 47.47.47.47/32 1 25600 25600 N/A xe6 12.1.1.4 Yes ELSP_CONFIG
R4:

```



```

R4#show ip ospf neighbor
Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID Pri State Dead Time Address Interface Instance ID
40.40.40.40 1 Full/Backup 00:00:33 13.1.1.1 xe2 0
58.58.58.58 1 Full/DR 00:00:39 14.1.1.3 xe3 0

R4#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
Transit RSVP:
To From Type LSPName State Uptime Rt Style Labelin Labelout DSType
47.47.47.47 40.40.40.40 PRI t2-Primary UP 00:54:24 1 1 SE 24320 52480 ELSP_CON
Total 1 displayed, Up 1, Down 0.
R4#show rsvp session detail
Transit
47.47.47.47
From: 40.40.40.40, LSPstate: Up, LSPname: t2-Primary
Transit upstream state: Operational, downstream state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 24320, Label out: 52480,
Tspec rate: 0k, Fspec rate: 0k
Tunnel Id: 5002, LSP Id: 2202, Ext-Tunnel Id: 40.40.40.40
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 14.1.1.3, xe3 Upstream: 13.1.1.1, xe2
Path refresh: 30 seconds (RR enabled) (due in 26617 seconds)
Path lifetime: 157 seconds (due in 143 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 36430 seconds)
Resv lifetime: 157 seconds (due in 143 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: Received Explicit Route Detail :
13.1.1.2/32 strict
14.1.1.3/32 strict
15.1.1.4/32 strict
Session Explicit Route Detail :
14.1.1.3/32 strict
15.1.1.4/32 strict
Record route:
-----
IP Address Label
-----
13.1.1.1
<self>
14.1.1.3
15.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS DSCP_value EXP_value
Current Error:
Code : None, Value : None
Originated Node : None, Recorded Time : N/A
Trunk Type: mpls R4#
R4#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
S - SNMP, L - LDP, R - RSVP, C - CRLDP
B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
P - SR Policy, U - unknown
Code FEC/VRF/L2CKT ILM-ID In-Label Out-Label In-Intf Out-Intf/VRF Nexthop LSP-Type
R> 47.47.47.47/32 1 24320 52480 N/A xe3 14.1.1.3 ELSP_CONFIG

```

```

R5#show ip ospf neighbor
Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID Pri State Dead Time Address Interface Instance ID
56.56.56.56 1 Full/Backup 00:00:39 14.1.1.2 xe42 0
47.47.47.47 1 Full/DR 00:00:38 15.1.1.4 xe39 0
R5#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
Transit RSVP:
To From Type LSPName State Uptime Rt Style Labelin Labelout DStype
47.47.47.47 40.40.40.40 PRI t2-Primary UP 00:59:48 1 1 SE 52480 25601 ELSP_CON
Total 1 displayed, Up 1, Down 0.
R5#show rsvp session detail
Transit
47.47.47.47
From: 40.40.40.40, LSPstate: Up, LSPname: t2-Primary
Transit upstream state: Operational, downstream state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 52480, Label out: 25601,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5002, LSP Id: 2202, Ext-Tunnel Id: 40.40.40.40
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 15.1.1.4, xe39 Upstream: 14.1.1.2, xe42
Path refresh: 30 seconds (RR enabled) (due in 26344 seconds)
Path lifetime: 157 seconds (due in 149 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 37267 seconds)
Resv lifetime: 157 seconds (due in 132 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: Received Explicit Route Detail :
14.1.1.3/32 strict
15.1.1.4/32 strict
Session Explicit Route Detail :
15.1.1.4/32 strict
Record route:
-----
IP Address Label
-----
13.1.1.1
14.1.1.2

<self>
15.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS DSCP_value EXP_value
Current Error:
Code : None, Value : None
Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
R5#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
S - SNMP, L - LDP, R - RSVP, C - CRLDP
B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
U - unknown
Code FEC/VRF/L2CKT ILM-ID In-Label Out-Label In-Intf Out-Intf/VRF Nexthop pri LSP-Type
R> 47.47.47.47/32 1 52480 25601 N/A xe39 15.1.1.4 Yes ELSP_CONFIG
R5#
R5#

```

R6:

```

R6#show ip ospf neighbor
Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID Pri State Dead Time Address Interface Instance ID
48.48.48.48 1 Full/DR 00:00:30 12.1.1.3 xe6 0
58.58.58.58 1 Full/Backup 00:00:33 15.1.1.3 xe5 0
R6#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
Egress RSVP:
To From Type LSPName State Uptime Rt Style Labelin Labelout DStype
47.47.47.47 40.40.40.40 PRI t1-Primary UP 01:05:33 1 1 SE 25600 - ELSP_CON
47.47.47.47 40.40.40.40 PRI t2-Primary UP 01:04:57 1 1 SE 25601 - ELSP_CON
Total 2 displayed, Up 2, Down 0.
R6#
R6#show rsvp session detail
Egress
47.47.47.47
From: 40.40.40.40, LSPstate: Up, LSPname: t1-Primary
Egress FSM state: Operational

Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 25600, Label out: -,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 40.40.40.40
Bind value: 0, Oper state: NA, Alloc mode: NA
Upstream: 12.1.1.3, xe6
Path lifetime: 157 seconds (due in 143 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 19100 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: Received Explicit Route Detail :
12.1.1.4/32 strict
Record route:
-----
IP Address Label
-----
10.10.1.1
11.1.1.2
12.1.1.3
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS DSCP_value EXP_value
Recorded Time : N/A
Current Error:
Code : None, Value : None
Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Egress
47.47.47.47
From: 40.40.40.40, LSPstate: Up, LSPname: t2-Primary
Egress FSM state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 25601, Label out: -,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5002, LSP Id: 2202, Ext-Tunnel Id: 40.40.40.40

```

```

Bind value: 0, Oper state: NA, Alloc mode: NA
Upstream: 15.1.1.3, xe5
Path lifetime: 157 seconds (due in 156 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 30939 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: Received Explicit Route Detail :
15.1.1.4/32 strict
Record route:
-----
IP Address Label
-----
13.1.1.1
14.1.1.2
15.1.1.3
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS DSCP_value EXP_value
Recorded Time : N/A
Current Error:

Code : None, Value : None
Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
R6#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
S - SNMP, L - LDP, R - RSVP, C - CRLDP
B - BGP , K - CLI , V - LDP_VC, I - IGP_SHORTCUT
O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
P - SR Policy, U - unknown
LDP ilm-ecmp - disabled
Code FEC/VRF/L2CKT ILM-ID In-Label Out-Label In-Intf Out-Intf/VRF Nexthop pri LSP-Type
R> 47.47.47.47/32 1 25600 Nolabel N/A N/A 127.0.0.1 Yes ELSP_CONFIG
R> 47.47.47.47/32 2 25601 Nolabel N/A N/A 127.0.0.1 Yes ELSP_CONFIG

```

RSVP-TE Dynamic Facility Backup LSP (RSVP Auto Bypass)

Overview

Resource Reservation Protocol (RSVP) auto bypass is a subset of the Facility Backup feature, designed to enhance fast-reroute protection. The facility backup feature operates by establishing bypass tunnels for protected sessions at each PLR node, serving as a local safeguard for sessions on every Point of Local Repair (PLR). Configuring bypass tunnels manually on each PLR, particularly in larger topologies, presented challenges in configuration management. The RSVP auto bypass functionality ensures to dynamically create bypass tunnels when feature enabled in PLR nodes and the sessions come up in PLR requesting facility backup protection.

Feature Characteristics

The RSVP auto bypass operates in conjunction with manual bypass tunnels, and if both manual and auto bypass offer the same level of protection as requested, the manual bypass takes precedence. Auto bypass tunnels, which do not serve any session, have a finite lifetime and are deleted if they remain unused.



Note: The auto bypass feature is not enabled by default and must be configured on every node of the network where auto bypass formation is required. However, even the nodes where auto-bypass is not enabled will continue to support transit or egress sessions of auto-bypass.

Benefits

RSVP auto bypass contributes to a more automated, efficient, and manageable facility backup protection mechanism in RSVP-enabled networks.

Prerequisites

Ensure the following prerequisites are met before configuring RSVP Auto Bypass functionality:

Define Interfaces and Loopback Addresses

Configure Layer 2 interfaces, such as port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification, and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 135.1.1.27/32 secondary  
  ipv6 address ::1/128  
!  
interface po6  
  ip address 1.1.2.2/30  
!  
interface xe6  
  channel-group 6 mode active  
!
```

Configure IGP for Dynamic Routing

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

ISIS Configuration

```
!  
router isis 1  
  is-type level-2-only  
  metric-style wide  
  mpls traffic-eng router-id 135.1.1.27  
  mpls traffic-eng level-2  
  capability cspf  
  dynamic-hostname  
  fast-reroute ti-lfa level-2 proto ipv4  
  net 49.0000.0000.0027.00  
  passive-interface lo  
!  
interface po6  
  isis network point-to-point  
  ip router isis 1  
!
```

OSPF Configuration

```
!  
router ospf 100  
  ospf router-id 135.1.1.27  
  network 135.1.1.27/32 area 0.0.0.0  
  network 1.1.5.1/24 area 0.0.0.0  
  network 1.1.1.1/24 area 0.0.0.0  
!
```

Configure RSVP for Efficient Network Operation

Enable Resource Reservation Protocol (RSVP) on all nodes to optimize traffic routing and quality of service. RSVP reserves network resources along specified paths to enhance network performance and reliability.

```
!  
router rsvp  
!  
interface xel  
  label-switching  
  enable-rsvp  
!
```

Configure the RSVP Primary Path and Trunk

Establish a trunk on edge routers participating in label-switching using defined path.

```
!  
rsvp-path PE1-PE4-1 mpls  
  1.1.1.2 strict  
  1.1.2.1 strict  
  1.1.6.1 strict
```

```

!
rsvp-trunk PE1-PE4-1 ipv4
primary fast-reroute protection facility
primary fast-reroute node-protection
primary path PE1-PE4-1
from 135.1.1.27
to 135.1.1.44
!

```

Configuration for RSVP Auto Bypass

This section shows the configuration of the various nodes within the topology to set up a RSVP-Auto bypass tunnels.

Topology

The sample topology includes Edge Nodes (PE1 and PE2) and core Nodes (P1, P3, and P4). As per sample configurations, Primary path is defined via PE1-P1-P3-PE2 using strict hops and auto bypass tunnel creation is formed on PE1.

Figure 30. RSVP-Auto Bypass Configuration



Before configuring RSVP-Auto bypass tunnels, meet all [Prerequisites \(page 721\)](#) for the following nodes:

- Edge nodes: PE1 and PE2
- Core nodes: P1, P3, and P4

Auto bypass feature helps in automatically forming bypass tunnels and protecting the facility backup protection requested tunnels.

Configuring RSVP Auto Bypass

Configure RSVP Auto bypass tunnels on PE1 Router

1. Configure auto-bypass under router rsvp.

```
(config)#router rsvp
(config-router)# auto-bypass
```

2. Enable Auto bypass globally.

```
(config-auto-bypass)# enable
```

3. Modify protection capability attribute for best-effort type bypasses.

```
(config-auto-bypass)# attributes best-effort
(config-auto-bypass-attr)#protection-capability node
```

4. Modify other auto bypass attributes for best-effort type bypasses.



Note: Auto bypass re-optimization uses the global command to set the re-optimization timer.

```
(config-auto-bypass-attr)# reoptimize
```

5. Configure inactivity timer globally for automatically created bypass tunnels.



Note: The default value for the inactivity-timer is 60 minutes. The minimum timer is 1 and maximum timer is 240.

```
(config-auto-bypass)# inactivity-timer 5
```

Configure RSVP Auto bypass tunnels on P1, P3, P4, PE2 Router

1. Configure auto-bypass under router RSVP.

```
(config)#router rsvp
(config-router)# auto-bypass
```

2. Enable Auto bypass globally.

```
(config-auto-bypass)# enable
```

3. Modify protection capability attribute for best-effort type bypasses.

```
(config-auto-bypass)# attributes best-effort
(config-auto-bypass-attr)#protection-capability link
```

4. Modify other auto bypass attributes for best-effort type bypasses.



Note: Auto bypass re-optimization shall use the global command to set the re-optimization timer.

```
(config-auto-bypass-attr)# reoptimize
```

5. Configure inactivity timer globally for automatically created bypass tunnels.



Note: The default value for the inactivity-timer is 60 minutes. The minimum timer is 1 and maximum timer is 240.

```
(config-auto-bypass)# inactivity-timer 5
```


Snippet Configuration on PE1 Router

```
!
router rsvp
  auto-bypass
  attributes best-effort
  protection-capability node
  reoptimize
  exit
  inactivity-timer 5
  enable
  exit
!
```

Snippet Configuration on P1, P3, and P4 Router

```
#show running-config
!
router rsvp
  auto-bypass
  attributes best-effort
  protection-capability link
  reoptimize
  exit
  enable
  exit
!
```

Validation

Verify all sessions including auto bypass sessions summary information:

Bypass name will have a character B or D. BL for Link protection and BN is for Node protection.

```
PE1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime     Rt   Style  Labelin  Labelout
135.1.1.43  135.1.1.27  5001    2203     BPS    BN-135.1.1.42-135.1.1.43-Bypass
UP 00:02:46 1 1 SE   -        25604
135.1.1.44  135.1.1.27  5002    2202     PRI    PE1-PE4-1-
Primary      UP 00:13:21 1 1 SE   -        25600
Total 2 displayed, Up 2, Down 0.

Transit RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime     Rt   Style  Labelin  Labelout
135.1.1.43  135.1.1.42  5001    2201     PRI    BL-1.1.2.1-135.1.1.43-
Bypass      UP 00:10:11 1 1 SE   25601  25600
135.1.1.42  135.1.1.43  5002    2202     PRI    BL-1.1.2.2-135.1.1.42-
Bypass      UP 00:10:00 1 1 SE   25602  3
Total 2 displayed, Up 2, Down 0.

Egress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime     Rt   Style  Labelin  Labelout
135.1.1.27  135.1.1.44  5001    2201     PRI    Pe4-Pe1-
Primary      UP 00:13:15 1 1 SE   25600  -
135.1.1.27  135.1.1.42  5002    2202     PRI    BL-1.1.1.1-135.1.1.27-
Bypass      UP 00:10:00 1 1 SE   3      -
```

```

PE2#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|------------|------------|--------|----------|----------|------------------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 135.1.1.27 | 135.1.1.44 | 5001 | 2201 | PRI | Pe4-Pe1- | |
| Primary | | UP | 00:16:18 | 1 1 SE | - | 25601 |
| 135.1.1.43 | 135.1.1.44 | 5002 | 2202 | BPS | BL-1.1.6.2-135.1.1.43- | |
| Bypass | UP | N/A | 0 0 SE | - | | |

Total 2 displayed, Up 1, Down 1.

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|------------|------------|--------|----------|----------|------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 135.1.1.44 | 135.1.1.27 | 5002 | 2202 | PRI | PE1-PE4-1- | |
| Primary | | UP | 00:16:24 | 1 1 SE | 25600 | - |

Total 1 displayed, Up 1, Down 0.

Verify all sessions including auto bypass sessions detailed information:

```

PE1#show rsvp session detail
Egress
135.1.1.27
  From: 135.1.1.44, LSPstate: Up, LSPname: Pe4-Pe1-Primary
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: facility
  Fast-Reroute bandwidth : 0
  Protection type desired: Link
  Fast-Reroute Hop limit: 255
  Fast-Reroute Setup priority: 7, Hold priority: 0
  Label in: 25600, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 135.1.1.44
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Upstream: 1.1.1.2, xel
  Path lifetime: 157 seconds (due in 140 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 24093 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Enabled
  Admin Groups: none
  Received Explicit Route Detail :
    1.1.1.1/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  1.1.6.1         25601
  1.1.2.1         25601
  1.1.1.2         25600
  <self>
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type:  ELSP_CONFIG
  CLASS      DSCP_value      EXP_value
Recorded Time : N/A
  Current Error:
    Code : None, Value : None
    Originated Node : None, Recorded Time : N/A
  Trunk Type: mpls
Ingress (Bypass)

```

```

135.1.1.43
From: 135.1.1.27, LSPstate: Up, LSPname: BN-135.1.1.42-135.1.1.43-Bypass
Ingress FSM state: Operational
Establishment Time: 0s 11ms
Setup priority: 7, Hold priority: 0, HOP limit: 255
Auto-bypass Info: Exclude-Node 135.1.1.42, Egress 135.1.1.43, Cspf ISIS
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
LSP Re-Optimization: Enabled, Re-Optimization Timer: 1 minutes, Cspf Client: ISIS
IGP-Shortcut: Disabled, LSP metric: 20
LSP Protection: None
Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out: 25604,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5001, LSP Id: 2203, Ext-Tunnel Id: 135.1.1.27
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 1.1.5.2, xe48
Path refresh: 30 seconds (RR enabled) (due in 29821 seconds)
Resv lifetime: 157 seconds (due in 147 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: none
Exclude Node: 135.1.1.42
Session Explicit Route Detail :
  1.1.5.2/32 strict
  1.1.4.1/32 strict
Record route:
-----
IP Address      Label
-----
<self>
1.1.5.2
1.1.4.1
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Total LSP protected : 1, Bandwidth in use : 0
Transit
135.1.1.43
From: 135.1.1.42, LSPstate: Up, LSPname: BL-1.1.2.1-135.1.1.43-Bypass
Transit upstream state: Operational, downstream state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 25601, Label out: 25600,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 135.1.1.42
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 1.1.5.2, xe48 Upstream: 1.1.1.2, xe1
Path refresh: 30 seconds (RR enabled) (due in 29384 seconds)
Path lifetime: 157 seconds (due in 140 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 32128 seconds)
Resv lifetime: 157 seconds (due in 147 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Received Explicit Route Detail :
  1.1.1.1/32 strict

```

```

1.1.5.2/32 strict
1.1.4.1/32 strict
Session Explicit Route Detail :
1.1.5.2/32 strict
1.1.4.1/32 strict
Record route:
-----
IP Address      Label
-----
1.1.1.2
<self>
1.1.5.2
1.1.4.1
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Egress
135.1.1.27
From: 135.1.1.42, LSPstate: Up, LSPname: BL-1.1.1.1-135.1.1.27-Bypass
Egress FSM state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in:      3, Label out: -,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5002, LSP Id: 2202, Ext-Tunnel Id: 135.1.1.42
Bind value: 0, Oper state: NA, Alloc mode: NA
Upstream: 1.1.5.2, xe48
Path lifetime: 157 seconds (due in 147 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 21158 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Received Explicit Route Detail :
1.1.5.1/32 strict
Record route:
-----
IP Address      Label
-----
1.1.3.2
1.1.5.2
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Transit
135.1.1.42
From: 135.1.1.43, LSPstate: Up, LSPname: BL-1.1.2.2-135.1.1.42-Bypass
Transit upstream state: Operational, downstream state: Operational
Setup priority: 2, Hold priority: 2
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in:      25602, Label out:      3,
Tspec rate: 0, Fspec rate: 0

```

```

Tunnel Id: 5002, LSP Id: 2202, Ext-Tunnel Id: 135.1.1.43
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 1.1.1.2, xe1 Upstream: 1.1.5.2, xe48
Path refresh: 30 seconds (RR enabled) (due in 29400 seconds)
Path lifetime: 157 seconds (due in 147 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 40649 seconds)
Resv lifetime: 157 seconds (due in 140 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Received Explicit Route Detail :
  1.1.5.1/32 strict
  1.1.1.2/32 strict
Session Explicit Route Detail :
  1.1.1.2/32 strict
Record route:
-----
IP Address      Label
-----
1.1.4.1
1.1.5.2
<self>
1.1.1.2
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS DSCP_value EXP_value
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Ingress (Primary)
135.1.1.44
From: 135.1.1.27, LSPstate: Up, LSPname: PE1-PE4-1-Primary
Ingress FSM state: Operational
Establishment Time: 0s 16ms
Setup priority: 7, Hold priority: 0, HOP limit: 255
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: ISIS
IGP-Shortcut: Disabled, LSP metric: 30
LSP Protection: facility
Fast-Reroute bandwidth : 0
Protection type desired: Node
Fast-Reroute Hop limit: 255
Fast-Reroute Setup priority: 7, Hold priority: 0
Bypass trunk: BN-135.1.1.42-135.1.1.43, Merge Point: 1.1.2.1, MP Label: 25600
  Bypass OutLabel: 25604, OutIntf: xe48
  Protection provided -> Type: Node,BW: Best-effort, Match-Code: Perfect-match
Label in: -, Label out: 25600,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5002, LSP Id: 2202, Ext-Tunnel Id: 135.1.1.27
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 1.1.1.2, xe1
Path refresh: 30 seconds (RR enabled) (due in 29192 seconds)
Resv refresh: 0 seconds (due in 1 seconds)
Resv lifetime: 157 seconds (due in 140 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Enabled
Admin Groups: none
Configured Path: PE1-PE4-1 (in use)
Configured Explicit Route Detail :
  1.1.1.2/32 strict
  1.1.2.1/32 strict
  1.1.6.1/32 strict

```

```

Session Explicit Route Detail :
 1.1.1.2/32 strict
 1.1.2.1/32 strict
 1.1.6.1/32 strict
Record route:
LP = 1 -> PLR's Downstream link is protected      PU = 1 -> Protection is in use on PLR
NP = 1 -> PLR's Downstream neighbor is protected  BP = 1 -> BW protection available at PLR
-----
IP Address      Label      (LP, PU, NP, BP)
-----
<self>
1.1.1.2         25600      ( 1,  0,  0,  0)
1.1.2.1         25600      ( 0,  0,  0,  0)
1.1.6.1         25600      ( 0,  0,  0,  0)
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls

```

Verify particular Bypass session detail information:

```

PE1#show rsvp bypass BN-135.1.1.42-135.1.1.43
Ingress (Bypass)
135.1.1.43
  From: 135.1.1.27, LSPstate: Up, LSPname: BN-135.1.1.42-135.1.1.43-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 11ms
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  Auto-bypass Info: Exclude-Node 135.1.1.42, Egress 135.1.1.43, Cspf ISIS
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 1 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out: 25604,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2203, Ext-Tunnel Id: 135.1.1.27
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 1.1.5.2, xe48
  Path refresh: 30 seconds (RR enabled) (due in 29807 seconds)
  Resv lifetime: 157 seconds (due in 157 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: none
  Exclude Node: 135.1.1.42
  Session Explicit Route Detail :
    1.1.5.2/32 strict
    1.1.4.1/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  1.1.5.2
  1.1.4.1
  Style: Shared Explicit Filter
  Traffic type: controlled-load

```

```

Minimum Path MTU: 1500
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Total LSP protected : 1, Bandwidth in use : 0

```

Verify using below show command for LSP's Protected by the Bypass tunnel:

```

PE1#show rsvp bypass BN-135.1.1.42-135.1.1.43 protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
           3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: BN-135.1.1.42-135.1.1.43
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id Lsp-Id   Lsp-Name           Role           Ext_tnl_
id         Ingress  Egress             Match-Code Mapped-BW
5002      2202     PE1-PE4-1-        135.1.1.27    135.1.1.27    135.1.1.44    0    NA
Primary   Ingress
Total LSP protected : 1
Bandwidth in use : 0

```

Verify using below show command for mpls forwarding-table information:

```

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled
Code  FEC           FTN-ID  Nhlfe-ID  Tunnel-id  Pri  LSP-Type  Out-Label  Out-
Intf  ELC  Nexthop
R(b)>  135.1.1.43/32  15      37        5001      Yes  LSP_
DEFAULT 25604      xe48      No        1.1.5.2
R(t)>  135.1.1.44/32  9        19        5002      Yes  LSP_
DEFAULT 25600      xe1       No        1.1.1.2
R(t)>  135.1.1.44/32  14      38        5002      No   LSP_
DEFAULT 25600      xe48      No        135.1.1.43

1.1.5.2, label 25604)
(via

```

Shutdown the primary path and verify for node protection functionality:

Shutdown interface between PE1-P1, so that primary path goes Down and backup in use.

```

config)#interface xe1
(config-if)#shutdown
(config-if)#commit

```

Verify auto bypass sessions summary information:

```

PE1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To           From           Tun-ID  LSP-ID  Type  LSPName           State

```

```

Uptime   Rt   Style  Labelin  Labelout
135.1.1.43      135.1.1.27      5001      2203      BPS      BN-135.1.1.42-135.1.1.43-Bypass
UP    00:06:18  1 3 SE  -      25604
135.1.1.44      135.1.1.27      5002      2202      PRI      PE1-PE4-1-
Primary      BU    00:16:53  1 3 SE  -      25600
135.1.1.44      135.1.1.27      5002      2201      PRI      PE1-PE4-1-
Primary      DN    N/A      0 0 SE  -      -
Total 3 displayed, Up 2, Down 1.

Egress RSVP:
To      From      Tun-ID   LSP-ID   Type   LSPName      State
Uptime   Rt   Style  Labelin  Labelout
135.1.1.27      135.1.1.44      5001      2201      PRI      Pe4-Pe1-
Primary      UP    00:16:47  1 1 SE  25600      -
135.1.1.27      135.1.1.42      5002      2202      PRI      BL-1.1.1.1-135.1.1.27-
Bypass      UP    00:13:32  1 1 SE  3      -
Total 2 displayed, Up 2, Down 0

```

Verify particular Bypass session detail information:

```

PE1#show rsvp bypass BN-135.1.1.42-135.1.1.43
Ingress (Bypass)
135.1.1.43
  From: 135.1.1.27, LSPstate: Up, LSPname: BN-135.1.1.42-135.1.1.43-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 11ms
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  Auto-bypass Info: Exclude-Node 135.1.1.42, Egress 135.1.1.43, Cspf ISIS
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 1 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out: 25604,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2203, Ext-Tunnel Id: 135.1.1.27
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 1.1.5.2, xe48
  Path refresh: 30 seconds (RR enabled) (due in 29602 seconds)
  Resv lifetime: 157 seconds (due in 155 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: none
  Exclude Node: 135.1.1.42
  Session Explicit Route Detail :
    1.1.5.2/32 strict
    1.1.4.1/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  1.1.5.2
  1.1.4.1
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  Recorded Time : N/A
  Current Error:
    Code : None, Value : None
    Originated Node : None, Recorded Time : N/A
  Last Signaled Error:
    Code : None, Value : None
    Originated Node : None, Recorded Time : N/A
  Trunk Type: mpls
  Total LSP protected : 1, Bandwidth in use : 0

```


Verify using below show command for MPLS forwarding-table information:

```
PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled
Code   FEC                FTN-ID   Nhlfe-ID  Tunnel-id  Pri   LSP-Type   Out-Label   Out-
Intf   ELC   Nexthop
R(b)>   135.1.1.43/32    15      37        5001      Yes   LSP_
DEFAULT 25604      xe48     No        1.1.5.2
R(t)!   135.1.1.44/32      9       19        5002      Yes   LSP_
DEFAULT 25600      xe1      No        1.1.1.2
R(t)>   135.1.1.44/32    14      38        5002      No    LSP_
DEFAULT 25600      xe48     No        135.1.1.43
                                           (via
1.1.5.2, label 25604)
```

Verify using below show command for LSP's Protected by the Bypass tunnel:

```
PE1#show rsvp bypass BN-135.1.1.42-135.1.1.43 protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
           3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: BN-135.1.1.42-135.1.1.43
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id  Lsp-Id   Lsp-Name           Role      Ext_tnl_
id         Ingress  Egress             Match-Code Mapped-BW
5002      2202     PE1-PE4-1-        135.1.1.27 135.1.1.27 135.1.1.44 0      NA
Primary
Total LSP protected : 1
Bandwidth in use : 0
```

Auto Bypass and Manual Bypass Co-exists with Perfect Protection

Manual bypass and auto bypass coexist and when manual bypass already providing perfect protection, then manual bypass preferred over auto bypass. auto bypass won't be triggered.

Auto bypass is created when Manual bypass is provided with perfect protection as link

Snippet Configuration on PE1 Router for Auto-bypass

```
#show running-config
!
router rsvp
  auto-bypass
  attributes best-effort
  protection-capability link
  reoptimize
  exit
  inactivity-timer 5
  enable
  exit
!
```

Snippet Configuration on PE1 Router for Manual-bypass

```
!
rsvp-path backup-PE1-P4-P1 mpls
```

```

1.1.5.2 strict
1.1.3.2 strict
!
rsvp-bypass backup-to-p1
path backup-PE1-P4-P1
to 135.1.1.44
!

```

Validation

Verify all sessions including auto bypass sessions summary information:

Manual bypass and auto bypass coexist and when manual bypass already providing perfect protection, auto bypass session goes down when manual bypass exists.

Verify all RSVP sessions information: The output shows only manual bypass session information when Manual bypass and auto bypass coexist and when manual bypass already providing perfect protection.

```

PE1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
135.1.1.44  135.1.1.27  5001     2205     PRI     PE1-PE4-         -
Primary     UP    00:16:19 1 1 SE    -        25604
135.1.1.44  135.1.1.27  5002     2206     PRI     PE1-PE4-1-       -
Primary     UP    00:16:21 1 1 SE    -        25602
135.1.1.44  135.1.1.27  5003     2203     BPS     backup-to-p1-    -
Bypass     UP    00:05:18 1 1 SE    -        25601
Total 3 displayed, Up 3, Down 0.

Transit RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
135.1.1.28  135.1.1.42  5002     2202     PRI     BL-1.1.9.1-135.1.1.28-
Bypass     UP    00:16:19 1 1 SE    25601    3
Total 1 displayed, Up 1, Down 0.

Egress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
135.1.1.27  135.1.1.42  5001     2201     PRI     BL-1.1.1.1-135.1.1.27-
Bypass     UP    00:16:37 1 1 SE    3        -
135.1.1.27  135.1.1.44  5001     2201     PRI     Pe4-Pe1-         -
Primary     UP    00:16:37 1 1 SE    25600    -
Total 2 displayed, Up 2, Down 0.

```

Verify using below show command for LSP's Protected by the Bypass tunnel:

```

PE1#show rsvp bypass BN-135.1.1.42-135.1.1.43 protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
          3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: BN-135.1.1.42-135.1.1.43
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id  Lsp-Id  Lsp-Name          Role          Ext_tnl_
id         Ingress Egress            Match-Code    Mapped-BW
5001       2205    PE1-PE4-1-        135.1.1.27    135.1.1.27    135.1.1.44    2        NA
Primary    Ingress
5002       2206    PE1-PE4-1-        135.1.1.27    135.1.1.27    135.1.1.44    2        NA
Primary    Ingress
Total LSP protected : 1

```

Bandwidth in use : 0

Auto Bypass with Node Protection and Manual Bypass with Link Protection

Snippet Configuration on PE1 Router for Auto-bypass

```
!
router rsvp
  auto-bypass
    attributes best-effort
    protection-capability node
    reoptimize
  exit
  inactivity-timer 5
  enable
  exit
!
```

Snippet Configuration on PE1 Router for Manual-bypass

```
!
rsvp-path backup-PE1-P4-P1 mpls
  1.1.5.2 strict
  1.1.3.2 strict
!
rsvp-bypass backup-to-p1
  path backup-PE1-P4-P1
  to 135.1.1.44
!
```

Validation

Verify all sessions including auto bypass sessions summary information: The output shows that Manual bypass for link protection is UP and auto bypass for node protection is UP and running.

```
PE1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime     Rt   Style  Labelin  Labelout
135.1.1.44  135.1.1.27  5001    2205     PRI    PE1-PE4-         -
Primary     UP    00:22:44 1 1 SE    25604
135.1.1.44  135.1.1.27  5002    2206     PRI    PE1-PE4-1-       -
Primary     UP    00:22:46 1 1 SE    25602
135.1.1.44  135.1.1.27  5003    2203     BPS    backup-to-p1-    25601
Bypass      UP    00:11:43 1 1 SE    -
135.1.1.28  135.1.1.27  5004    2201     BPS    BN-135.1.1.42-135.1.1.28-Bypass
UP          00:05:17 1 1 SE    -        3
135.1.1.43  135.1.1.27  5006    2204     BPS    BN-135.1.1.42-135.1.1.43-Bypass
UP          00:05:17 1 1 SE    -        25602
Total 5 displayed, Up 5, Down 0.

Transit RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime     Rt   Style  Labelin  Labelout
135.1.1.28  135.1.1.42  5002    2202     PRI    BL-1.1.9.1-135.1.1.28-
Bypass      UP    00:22:44 1 1 SE    25601    3
Total 1 displayed, Up 1, Down 0.

Egress RSVP:
```

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|----------------------------------|------------|-----------------|---------|----------|------------------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 135.1.1.27 | 135.1.1.42 | 5001 | 2201 | PRI | BL-1.1.1.1-135.1.1.27- | |
| Bypass | UP | 00:23:02 1 1 SE | 3 | - | | |
| 135.1.1.27 | 135.1.1.44 | 5001 | 2201 | PRI | Pe4-Pe1- | |
| Primary | UP | 00:23:02 1 1 SE | 25600 | - | | |
| Total 2 displayed, Up 2, Down 0. | | | | | | |

Verify using below show command for LSP's Protected by the Bypass tunnels:

```
PE1#sh rsvp bypass protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
           3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: backup-to-p1
Bypass trunk bandwidth type: best-effort
Total LSP protected : 0
Bandwidth in use : 0

Bypass trunk: BN-135.1.1.42-135.1.1.43
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id Lsp-Id Lsp-Name Role Ext_tnl_
id Ingress Egress Match-Code Mapped-BW
5006 2207 PE1-PE4-1- Primary Ingress 135.1.1.27 135.1.1.27 135.1.1.44 0 NA
Total LSP protected : 1
Bandwidth in use : 0

Bypass trunk: BN-135.1.1.42-135.1.1.28
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id Lsp-Id Lsp-Name Role Ext_tnl_
id Ingress Egress Match-Code Mapped-BW
5004 2205 PE1-PE4- Primary Ingress 135.1.1.27 135.1.1.27 135.1.1.44 0 NA
Total LSP protected : 1
Bandwidth in use : 0
```

RSVP Auto Bypass Commands

The RSVP auto bypass introduces the following configuration commands within the Facility Backup commands.

auto-bypass

Use this command to enter a new mode of auto-bypass to configure auto bypass specific commands and enable auto bypass feature.

Use the `no` parameter to auto bypass configurations. When configuration removed at global level, all auto bypass configuration will be removed and all dynamically created bypass tunnels will be deleted.

Command Syntax

```
auto-bypass
```

```
no auto-bypass
```

Parameters

None

Default

Disabled

Command Mode

Router mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

This example is for configuring auto bypass:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#exit
(config-router)#no auto-bypass
```

auto-bypass enable

Use this command to enable or disable auto bypass functionality. When this functionality is enabled, sessions requesting facility backup protection will attempt creation of bypass tunnels dynamically, if no satisfying bypass protection is available. If topology supports the formation of dynamically created bypass session, then the sessions will be protected by dynamically created bypass tunnels.

When manual bypass created, which satisfies protection requirement similar to auto bypass tunnels, then manual bypass tunnel gets preference over auto bypass tunnels. However, if auto bypass tunnels provide better protection (based on session request) than manual bypass, then auto bypass tunnels considered for protection over manual bypass tunnels.

Use the `disable` parameter to disable the auto bypass functionality. When this functionality is disabled, all dynamically created bypass tunnels will be deleted.

Command Syntax

```
enable
```

```
disable
```

Parameters

None

Default

Disabled

Command Mode

Auto bypass mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

This example is for enabling the auto bypass:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#enable
(config-auto-bypass)#commit
(config-auto-bypass)#exit
```

This example is for disabling the auto bypass:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#disable
(config-auto-bypass)#commit
(config-auto-bypass)#exit
```

inactivity-timer

Use this command to configure inactivity timer for auto bypass tunnels. When auto bypass tunnels do not protect any session, inactivity timer will decide how long to retain the auto bypass tunnels before deleting them. If auto bypass tunnel starts protecting any session, then inactivity timer will be stopped.

Use the `no` parameter to reset auto bypass inactivity timer to default value.

Command Syntax

```
inactivity-timer <1-240>
```

```
no inactivity-timer
```

Parameters

inactivity-timer <1-240>

Specifies the RSVP auto bypass inactivity timer in minutes. This timer determines the maximum time that dynamically created auto bypass tunnels can remain inactive before being deleted. The default value is 60 minutes.

Default

None

Command Mode

Auto bypass mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

This example is for setting the inactivity timer to 5 minutes:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#inactivity-timer 5
(config-auto-bypass)#commit
(config-auto-bypass)#exit
```


attributes

Use this command to configure auto bypass attributes. Currently only best-effort type auto bypass tunnels supported and hence user can modify the attributes of best-effort type auto bypass. This command takes to new mode in which several parameters can be configured. The parameters configured in this mode are applicable to auto bypass sessions of attribute type.

Use the `no` parameter to reset auto bypass attributes of matching type to default values.

Command Syntax

```
attributes best-effort
```

```
no attributes best-effort
```

Parameters

best-effort

Specifies the auto bypass attribute configuration for the best-effort type in auto bypass.

Default

The initial configuration for best-effort auto bypass attributes is not set, and the internal entities have their default settings.

Command Mode

Auto bypass mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

This example is for configuring auto bypass attributes:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

This example is to reset auto bypass attributes of matching type to default values:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
(config-auto-bypass)#no attributes best-effort
(config-auto-bypass)#commit
(config-auto-bypass)#exit
```

protection-capability

Use this command to set the protection capability attribute for auto bypass. Capability setting decides the types of auto bypass tunnels allowed to be attempted. When capability configured as link, only exclude link type of bypass creation will be attempted and provides protection against only link failure. When capability node is configured, if session requests for node protection, then exclude node type bypass will also be created along with exclude link type bypass and provides the best protection possible. Even when capability configured as node, if session doesn't request for node protection, then node protecting bypass formation will not be attempted.

Use the `no` parameter to reset the attribute to default value which is capability type node.

Command Syntax

```
protection-capability (link | node)
```

```
no protection-capability
```

Parameters

link

Specifies the capability for link protection only.

node

Specifies the capability for link and node protection.

Default

The value for protection capability is `node`.

Command Mode

Auto bypass attribute mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for setting up the protection capability attribute to `link`:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#protection-capability link
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

Example to reset the protection capability attribute to default value:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#no protection-capability
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

bandwidth

Use this command to set the bandwidth attribute for auto bypass. Auto bypass sessions of matching attribute type will inherit this bandwidth constraint while requesting for an LSP to CSPF.

Use the `no` parameter to remove bandwidth attribute and auto bypass sessions will not request any bandwidth constraint to form LSP.

Command Syntax

```
bandwidth BANDWIDTH
```

```
no bandwidth
```

Parameters

BANDWIDTH

Specifies the bandwidth attribute for auto bypass within the range of 1k to 100g. The default value is zero.

Default

None

Command Mode

Auto bypass attribute mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for setting up the bandwidth attribute to 100g:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#bandwidth 100g
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

Example to reset the bandwidth attribute to default value:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#no bandwidth
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

setup-priority

Use this command to set the setup priority attribute for auto bypass. Auto bypass sessions of matching attribute type will inherit this setup priority constraint while requesting for an LSP to CSPF.

Use the `no` parameter to reset setup priority attribute to default value.

Command Syntax

```
setup-priority <0-7>
```

```
no setup-priority
```

Parameters

setup-priority <0-7>

Specifies the attribute for Auto bypass setup priority. The default value for attribute for setup priority is 7.

Default

None

Command Mode

Auto bypass attribute mode

Applicability

Introduced in OcNOS version 6.5.0.

Example

Example for setting up the setup priority attribute to 5.

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass attributes best-effort
(config-auto-bypass-attr)#setup-priority 5
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

Example to reset the setup priority attribute to default value.

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass attributes best-effort
(config-auto-bypass-attr)#no setup-priority
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

hold-priority

Use this command to set the hold priority attribute for auto bypass. Auto bypass sessions of matching attribute type will inherit this hold priority constraint while requesting for an LSP to CSPF.

Use the `no` parameter to reset hold priority attribute to default value.

Command Syntax

```
hold-priority <0-7>
```

```
no hold-priority
```

Parameters

hold-priority <0-7>

Specifies the attribute for Auto bypass hold priority. The default value is 0.

Default

None

Command Mode

Auto bypass attribute mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for setting up the hold priority attribute to 2:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#hold-priority 2
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

Example to reset the hold priority attribute to default value:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#no hold-priority
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

hop-limit

Use this command to set the hop limit attribute for auto bypass. Auto bypass sessions of matching attribute type will inherit this hop limit constraint while requesting for an LSP to CSPF.

Use the `no` parameter to reset hop limit attribute to default value.

Command Syntax

```
hop-limit <1-255>
```

```
no hop-limit
```

Parameters

hop-limit <1-255>

Specifies the attribute for Auto bypass hop limit. The default value is 255.

Default

None

Command Mode

Auto bypass attribute mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for setting up the hop limit attribute to 5:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#hop-limit 5
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

Example to reset the hop limit attribute to default value:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#no hop-limit
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

reoptimize

Use this command to set the re-optimization attribute for auto bypass. Auto bypass sessions of matching attribute type will be enabled for re-optimization when the attribute is configured.

Use the `no` parameter to disable re-optimization for auto bypass.

Command Syntax

```
reoptimize
```

```
no reoptimize
```

Parameters

None

Default

The auto bypass re-optimization is disabled.

Command Mode

Auto bypass attribute mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for enabling the re-optimization attribute:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#reoptimize
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

Example for disabling the re-optimization attribute:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#no reoptimize
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

exclude-address

Use this command to configure exclude address attribute for auto bypass. Auto bypass sessions of matching attribute type will inherit this explicit exclude address constraint while requesting for an LSP to CSPF. CSPF will provide the auto bypass path excluding the implicit exclude address and the configured exclude addresses.



Note: When link address added for exclusion, next-hop address on the link must be configured to exclude the link in forward direction and also if the link must be excluded in both direction, then both link address and the next-hop address must be added to exclude address list. For node address exclusion, it is recommended to configure the loopback address which is part of routing.

Use the `no` parameter to remove the exclude address constraint.

Command Syntax

```
exclude-address (link | node) A.B.C.D
```

```
no exclude-address (link | node) A.B.C.D
```

Parameters

link A.B.C.D

Specifies the exclude address of type link.

node A.B.C.D

Specifies the exclude address of node link.

Default

Auto bypass will only have an implicit exclude address, which is the immediate link or node of the protected session. No additional exclude address constraints are applied.

Command Mode

Auto bypass attribute mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for configuring exclude address attribute to link 10.1.1.2:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#exclude-address link 10.1.1.2
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

Example for configuring exclude address attribute to node 2.2.2.2:

```
#configure terminal
(config)#router rsvp
```



```
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#exclude-address node 2.2.2.2
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

Example to remove the exclude address constraint from link 10.1.1.2:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bypass
(config-auto-bypass)#attributes best-effort
(config-auto-bypass-attr)#no exclude-address link 10.1.1.2
(config-auto-bypass-attr)#commit
(config-auto-bypass-attr)#exit
```

show rsvp router-id-table

Use this command to display the mapping of link addresses to router IDs maintained by the RSVP module to reduce the number of CSPF query.



Note: This command can display no information even when auto bypass sessions are active but that is just an indication that RSVP module doesn't hold the CSPF data locally and will collect data from CSPF when required.

Command Syntax

```
show rsvp router-id-table
```

Parameters

None

Command Mode

Exec and Privileged exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for viewing RSVP router ID table:

```
#show rsvp router-id-table
Link Address      Router ID      Query Status   CSPF Client     Refcnt
-----+-----+-----+-----+-----
23.1.1.3          3.3.3.3        COMPLETE       ISIS             3
```

show rsvp auto-bypass-group

Use this command to display the mapping of the group ID value used in the auto bypass name to exclude interface index and SRLG values. This command helps understand the value present in the auto bypass tunnel names.

Command Syntax

```
show rsvp auto-bypass-group
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

Example for viewing all the auto bandwidth profiles:

```
#show rsvp auto-bypass-group
SRLG-disjoint Configured: Forced
Group-ID  Refcnt    Exclude-IfIndex  SRLG Values
-----+-----+-----+-----
100       3           10001           10 20
```


RSVP Graceful Restart Configuration

Overview

The RSVP-TE graceful restart enables routers to maintain MPLS forwarding state when a link or node failure occurs. In a link failure, control communication is lost between two nodes, however, the nodes do not lose their control or forwarding state. RSVP Graceful restart (GR) is one of the fault-handling mechanism, that protects the forwarding state of the node during failure and helps to reinstate the previous state when the node has recovered.

Topology

Figure 31. RSVP Graceful Restart



Configuration

PE1-NSM

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 11.11.11.11/32 secondary | Set the IP address for the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe7 | Enter interface mode. |
| (config-if)#ip address 10.10.10.2/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe7. |
| (config-if)#commit | Commit the transaction. |

PE1-RSVP-TE

| | |
|-------------------------|---|
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe7 | Enter interface mode. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#commit | Commit the transaction. |

PE1-OSPF

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the Routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 11.11.11.11 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 10.10.10.0/24 area 0 | Define the network (10.10.10.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 11.11.11.11/32 area 0 | Set the IP address of the loopback interface to 11.11.11.11/32. |
| (config-router)#commit | Commit the transaction. |

P - NSM

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address for the interface. |
| (config-if)#exit | Enable label switching on interface lo. |
| (config)#interface xe7 | Enter interface mode. |
| (config-if)#ip address 10.10.10.1/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe7. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe8 | Enter interface mode. |
| (config-if)#ip address 20.20.20.1/24 | Set the IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe7. |
| (config-if)#commit | Commit the transaction. |

P - RSVP-TE

| | |
|-------------------------|---|
| (config)#router rsvp | Enter Configure Router mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe7 | Enter interface mode. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe8 | Enter interface mode. |

| | |
|--------------------------------------|---|
| <code>(config-if)#enable-rsvp</code> | Enable RSVP message exchange on this interface. |
| <code>(config-if)#commit</code> | Commit the transaction. |

P - OSPF

| | |
|---|---|
| <code>#configure terminal</code> | Enter configure mode. |
| <code>(config)#router ospf 100</code> | Configure the Routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| <code>(config-router)#router-id 1.1.1.1</code> | Configure OSPF router-ID same as loopback interface IP address |
| <code>(config-router)#network 10.10.10.0/24 area 0</code> | Define the first network (10.10.10.0/24) on which OSPF runs and associate the area ID (0). |
| <code>(config-router)#network 20.20.20.0/24 area 0</code> | Define the second network (20.20.20.0/24) on which OSPF runs and associate the area ID (0). |
| <code>(config-router)#network 1.1.1.1/32 area 0</code> | Set the IP address of the loopback interface to 1.1.1.1/32. |
| <code>(config-router)#commit</code> | Commit the transaction. |

PE3 - NSM

| | |
|--|--|
| <code>#configure terminal</code> | Enter configure mode. |
| <code>(config)#interface lo</code> | Enter interface mode. |
| <code>(config-if)#ip address 33.33.33.33/32 secondary</code> | Set the IP address for the interface. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface ge8</code> | Enter interface mode. |
| <code>(config-if)#ip address 20.20.20.2/24</code> | Set the IP address for the interface. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface ge8. |
| <code>(config-if)#commit</code> | Commit the transaction. |

PE3- RSVP-TE

| | |
|--------------------------------------|---|
| <code>(config)#router rsvp</code> | Enter Configure Router mode. |
| <code>(config-router)#exit</code> | Exit Router mode. |
| <code>(config)#interface ge8</code> | Enter interface mode. |
| <code>(config-if)#enable-rsvp</code> | Enable RSVP message exchange on this interface. |
| <code>(config-if)#commit</code> | Commit the transaction. |

PE3 - OSPF

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the Routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#router-id 33.33.33.33 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 20.20.20.0/24 area 0 | Define the network (20.20.20.0/24) on which OSPF runs and associate the area ID (0). |
| (config-router)#network 33.33.33.33/32 area 0 | Set the IP address of the loopback interface to 33.33.33.33/32. |
| (config-router)#commit | Commit the transaction. |

PE1 - RSVP-Path

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#rsvp-path PE1_to_PE3_via_P1 mpls | Create an RSVP Path and enter the Path mode. |
| (config-path)#10.10.10.1 strict | Configure this explicit route path as a strict hop. |
| (config-path)#20.20.20.2 strict | Configure this explicit route path as a strict hop. |
| (config-path)#exit | Exit Path mode. |
| #configure terminal | Enter configure mode. |
| (config)#rsvp-trunk PE1_to_PE3 ipv4 | Create an RSVP trunk and enter the Trunk mode. |
| (config-trunk)#primary path PE1_to_PE3_via_P1 | Configure the trunk to use defined path. |
| (config-trunk)#from 11.11.11.11 | Specify ipv4 source point for the LSP. |
| (config-trunk)#to 33.33.33.33 | Specify the IPv4 egress (destination point) for the LSP. |
| (config-trunk)#commit | Commit the transaction. |

PE3 - RSVP-Path

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#rsvp-path PE3_to_PE1_via_P1 mpls | Create an RSVP Path and enter the Path mode. |
| (config-path)#20.20.20.1 strict | Configure this explicit route path as a strict hop. |
| (config-path)#10.10.10.2 strict | Configure this explicit route path as a strict hop. |
| (config-path)#exit | Exit Path mode. |
| #configure terminal | Enter configure mode. |
| (config)#rsvp-trunk PE3_to_PE1 ipv4 | Create an RSVP trunk and enter the Trunk mode. |
| (config-trunk)#primary path PE3_to_PE1_via_P1 | Configure the trunk to use defined path. |

| | |
|---------------------------------|--|
| (config-trunk)#from 33.33.33.33 | Specify the ipv4 source point for LSP |
| (config-trunk)#to 11.11.11.11 | Specify the IPv4 egress (destination point) for the LSP. |
| (config-trunk)#commit | Commit the transaction. |

PE1 – RSVP-GR

| | |
|-------------------------------------|--------------------------------------|
| #configure terminal | Enter configure mode. |
| (config)#router rsvp | Enter the configuration router mode. |
| (config-router)#neighbor 10.10.10.1 | Configure the neighbor ip address. |
| (config-router)#graceful-restart | Enable the rsvp-gr. |
| (config-router)#commit | Commit the transaction. |

P1 – RSVP-GR

| | |
|-------------------------------------|--------------------------------------|
| #configure terminal | Enter configure mode. |
| (config)#router rsvp | Enter the configuration router mode. |
| (config-router)#neighbor 10.10.10.2 | Configure the neighbor ip address. |
| (config-router)#neighbor 20.20.20.2 | Configure the neighbor ip address. |
| (config-router)#graceful-restart | Enable the rsvp-gr. |
| (config-router)#commit | Commit the transaction. |

PE3 – RSVP-GR

| | |
|-------------------------------------|--------------------------------------|
| #configure terminal | Enter configure mode. |
| (config)#router rsvp | Enter the router configuration mode. |
| (config-router)#neighbor 20.20.20.1 | Configure the neighbor ip address. |
| (config-router)#graceful-restart | Enable the rsvp-gr. |
| (config-router)#commit | Commit the transaction. |

Validation

Verify the RSVP graceful restart on PE1.

```

PE1#show rsvp graceful-restart
Graceful Restart: Enabled
Advertised Restart Time: 200000 msec
Advertised Recovery Time: 360000 msec
Sending Recovery Time: No

PE1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```



```
Egress RSVP:
To          From          Tun-ID  LSP-
ID  Type    LSPName      State  Uptime    Rt   Style  Labelin  Labelout
11.11.11.11 33.33.33.33 5001    2201    PRI    PE1_to_PE3 ipv4s   UP    00:54:57 1
1  SE      3          -
```

Total 5 displayed, Up 5, Down 0.

Verify after performing RSVP graceful restart on PE1

```
7038-PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled

Code      FEC              FTN-ID   Nhife-ID  Tunnel-id  Pri   LSP-Type   Out-Label  Out-
Intf      ELC    Nexthop
L>        1.1.1.1/32      3         4         -          Yes   LSP_
          3              xe7        No        10.10.10.1
DEFAULT  3              20         81        5001       Yes   LSP_
R(t)> p 33.33.33.33/32    xe7        Yes p 10.10.10.1
DEFAULT  26881          xe7
```

Fast Reroute Configuration (One-to-one Method)

The Fast Reroute (FRR) configuration is a MPLS resiliency technology that provides fast traffic recovery when there is a link or router failure on mission critical services. These mechanisms enable the re-direction of traffic onto backup LSP tunnels in tens of milliseconds, in the event of a failure. The one-to-one backup method creates detour LSPs for each protected LSP at each potential point of local repair. This method is used to protect links and nodes during network failure.

In the below configurations each FRR trunk is mapped to VPWS,VPLS, and L3 VPN services. So it includes configurations of VPWS,VPLS and L3 VPN also.

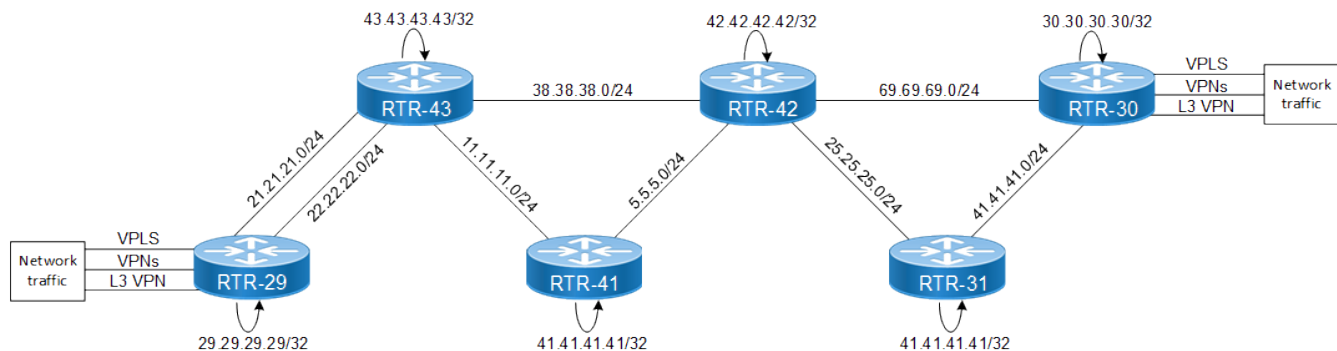


Note: IP BFD must be configured at 3.3ms to achieve 50ms switchover for RSVP FRR.



Note: Starting with OcNOS version 7.0.0, sub-interfaces replace Service Templates as the preferred configuration method. While legacy Service Template configurations remain functional, the associated commands are now hidden and will be completely removed in a future release. To ensure future-proofing and continued manageability, please migrate all existing Service Templates to sub-interfaces. For more details, see [Appendix A: Service Template Configuration \(page 2631\)](#).

Figure 32. Topology Example for Fast Reroute



RTR-29 Configuration

The following are the RTR-29 router configuration procedures.

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 29.29.29.29/32 secondary | Set a secondary IP address of the interface |
| (config-if)#no shutdown | Administratively bring the interface up. |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 29.29.29.29 | Configure the router ID. |

| | |
|---|---|
| (config)#router rsvp | Enter to router rsvp mode. |
| (config-router)#exit | Exit the router mode |
| (config)#router ldp | Enter to router LDP mode. |
| (config-router)#targeted-peer ipv4 30.30.30.30 | Configure targeted peer. |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 21.21.21.29/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on this interface |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe22 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on this interface |
| (config-if)#ip address 22.22.22.29/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf | Enter the router configure mode for OSPF. |
| (config-router)#router-id 29.29.29.29 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 21.21.21.0/24 area 0 (config-router)#network 22.22.22.0/24 area 0 (config-router)#network 29.29.29.29/32 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#exit | Exit the router configure mode. |
| (config)#rsvp-path p21 | Enter the path mode for RSVP pt1. |
| (config-path)#21.21.21.43 strict | Configure this explicit route path as a strict hop. |
| (config-path)#38.38.38.42 strict | Configure this explicit route path as a strict hop. |
| (config-path)#69.69.69.30 strict | Configure this explicit route path as a strict hop. |
| (config)#exit | Exit the path mode. |
| (config)#rsvp-trunk to_30 ipv4 | Enter the trunk mode for RSVP. |
| (config-trunk)#primary fast-reroute protection one-to-one | Configure primary fast-reroute protection facility for a trunk. |

| | |
|--|--|
| (config-trunk)#primary fast-reroute node-protection | Configure primary fast-reroute node protection for the trunk |
| (config-trunk)#primary path p21 | Configure trunk to 30 to use the defined path. |
| (config-trunk)#to 30.30.30.30 | Specify the IPv4 egress (destination point) for the LSP. |
| (config-trunk)#exit | Exit from trunk mode. |
| (config)#ip vrf vrf1 | Configure VRF instance |
| (config-vrf)#rd 100:1 | Configure Router Distinguisher value |
| (config-vrf)#route-target both 100:1 | Configure route-target as both |
| (config-vrf)#exit | Exit the path mode. |
| (config)#interface xe43 | Enter to the interface mode |
| (config-if)#ip vrf forwarding vrf1 | Bind the VRF instance to the interface |
| (config-if)#ip address 43.43.43.29/24 | Configure IP address |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 100 | Configure BGP router instance |
| (config-router)#neighbor 30.30.30.30 remote-as 100 | Configure neighbor with remote-as |
| (config-router)#neighbor 30.30.30.30 update-source 29.29.29.29 | Configure update source as loopback address |
| (config-router)#address-family vpnv4 unicast | Configure VPNv4 address family |
| (config-router-af)#neighbor 30.30.30.30 activate | Activate the VPN neighbor |
| (config-router-af)#exit-address-family | Exit the VPN address family |
| (config-router)#address-family ipv4 vrf vrf1 | Configure VRF address family |
| (config-router-af)redistribute connected | Redistribute connected route |
| (config-router-af)exit-address-family | Exit VRF address family |
| (config-router)#exit | Exit router mode |
| (config)#mpls l2-circuit vlan10 10 30.30.30.30 | Configure Virtual circuit. |
| (config-pseudowire)#exit | Exit pseudowire config mode. |
| (config)#service-template st1 | Template configuration |
| (config-svc)#match outer-vlan 10 | Match criteria under template configuration |
| (config-svc)#exit | Exit service template mode |
| (config)#service-template st2 | Template configuration |
| (config-svc)#match outer-vlan 30 | Match criteria under template configuration |
| (config-svc)#exit | Exit service template mode |
| (config)#interface xe44 | Enter interface configuration mode |
| (config-if)#switchport | Configure interface as switch port |
| (config-if)#mpls-l2-circuit t1 service-template st1 | Bind the interface to the VC with service template |

| | |
|---|---|
| (config-if)#exit | Exit interface configuration mode |
| (config)#mpls vpls vpls30 30 | Configure VPLS instance |
| (config-vpls)#signaling ldp | Configure VPLS signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Configure VPLS type as VLAN encapsulation |
| (config-vpls-sig)#vpls-peer 30.30.30.30 | Configure VPLS peer |
| (config-vpls-sig)#exit-signaling | Exit VPLS configuration mode |
| (config)#interface xe45 | Enter interface configuration mode |
| (config-if)#switchport | Configure interface as switch port |
| (config-if)#mpls-vpls vpls30 service-template st2 | Bind the VPLS instance to the interface |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config-if)#commit | Commit the transaction. |

RTR-43 Configuration

The following are the RTR-43 router configuration procedures.

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 43.43.43.43/32 secondary | Set a secondary IP address of the interface |
| (config-if)#no shutdown | Administratively shutdown the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 43.43.43.43 | Configure the router ID. |
| (config)#router rsvp | Enter to router RSVP mode. |
| (config-router)#exit | Exit the router mode |
| (config)#interface xe5/1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 11.11.11.43/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe9/1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface |
| (config-if)#ip address 21.21.21.43/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this |

| | |
|---|---|
| | interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe9/2 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 22.22.22.43/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe13/2 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface |
| (config-if)#ip address 38.38.38.43/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf | Enter the router configure mode for OSPF. |
| (config-router)#router-id 43.43.43.43 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 11.11.11.0/24 area 0 (config-router)#network 22.22.22.0/24 area 0 (config-router)#network 21.21.21.0/24 area 0 (config-router)#network 38.38.38.0/24 area 0 (config-router)#network 43.43.43.43/32 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#commit | Commit the transaction. |

RTR-42 Configuration

The following are the RTR-42 router configuration procedures.

| | |
|---|---|
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 42.42.42.42/32 secondary | Set a secondary IP address of the interface |
| (config-if)#no shutdown | Administratively shutdown the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 42.42.42.42 | Configure the router ID. |
| (config)#router rsvp | Enter to router RSVP mode. |
| (config-router)#exit | Exit the router mode |
| (config)#interface xe2 | Enter interface mode. |

| | |
|--|---|
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 5.5.5.42/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe10/1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface |
| (config-if)#ip address 25.25.25.42/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe3 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 38.38.38.42/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe4 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface |
| (config-if)#ip address 69.69.69.42/24 | Specify an IP address for the interface |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf | Enter the router configure mode for OSPF. |
| (config-router)#router-id 42.42.42.42 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 5.5.5.0/24 area 0 (config-router)#network 25.25.25.0/24 area 0 (config-router)#network 69.69.69.0/24 area 0 (config-router)#network 38.38.38.0/24 area 0 (config-router)#network 42.42.42.42/32 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#commit | Commit the transaction |

RTR-41 Configuration

The following are the RTR-41 router configuration procedures.

| | |
|---|---|
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 44.44.44.44/32 secondary | Set a secondary IP address of the interface |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 44.44.44.44 | Configure the router ID. |
| (config)#router rsvp | Enter to router RSVP mode. |
| (config-router)#exit | Exit the router mode |
| (config)#interface xe1/1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 1.1.1.41/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface |
| (config-if)#ip address 5.5.5.41/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe5/1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 11.11.11.41/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf | Enter the router configure mode for OSPF. |
| (config-router)#router-id 44.44.44.44 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 5.5.5.0/24 area 0 (config-router)#network 1.1.1.0/24 area 0 (config-router)#network 11.11.11.0/24 area 0 (config-router)#network 44.44.44.44/32 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#commit | Commit the transaction. |

RTR-31 Configuration

The following are the RTR-31 router configuration procedures.

| | |
|--|---|
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 31.31.31.31/32 secondary | Set a secondary IP address of the interface |
| (config-if)#no shutdown | Administratively shutdown the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 31.31.31.31 | Configure the router ID. |
| (config)#router rsvp | Enter to router RSVP mode. |
| (config-router)#exit | Exit the router mode |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 1.1.1.31/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe25 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface |
| (config-if)#ip address 25.25.25.31/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe41 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 41.41.41.31/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf | Enter the router configure mode for OSPF. |
| (config-router)#router-id 31.31.31.31 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 1.1.1.0/24 area 0 (config-router)#network 25.25.25.0/24 area 0 (config-router)#network 41.41.41.0/24 area 0 (config-router)#network 31.31.31.31/32 area 0 | Define the network on which OSPF runs and associate the area ID |
| (config-router)#commit | Commit the transaction. |

RTR-30 Configuration

The following are the RTR-30 router configuration procedures.

| | |
|---|---|
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 30.30.30.30/32 secondary | Set a secondary IP address of the interface |
| (config-if)#no shutdown | Administratively shutdown the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router-id 30.30.30.30 | Configure the router ID. |
| (config)#router rsvp | Enter to router RSVP mode. |
| (config-router)#exit | Exit the router mode |
| (config)#router ldp | Enter to router LDP mode. |
| (config-router)#targeted-peer ipv4 29.29.29.29 | Configure targeted peer. |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |
| (config)#interface xe41 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface. |
| (config-if)#ip address 41.41.41.30/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on this interface |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe54/1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching on interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on this interface |
| (config-if)#ip address 69.69.69.30/24 | Set an IP address of the interface. |
| (config-if)#no shutdown | Administratively no shutdown the interface. |
| (config-if)#enable-rsvp | Enable RSVP message exchange on this interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf | Enter the router configure mode for OSPF. |
| (config-router)#router-id 30.30.30.30 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router)#network 41.41.41.0/24 area 0 (config-router)#network 69.69.69.0/24 area 0 (config-router)#network 30.30.30.30/32 area 0 | Define the network on which OSPF runs and associate the area ID |

| | |
|--|---|
| (config-router)#commit | Commit the transaction. |
| (config)#rsvp-path p41 | Enter the path mode for RSVP pt1. |
| (config-path)#41.41.41.31 strict | Configure this explicit route path as a strict hop. |
| (config-path)#1.1.1.41 strict | Configure this explicit route path as a strict hop. |
| (config-path)#11.11.11.43 strict | Configure this explicit route path as a strict hop. |
| (config)#exit | Exit the path mode. |
| (config)#rsvp-trunk to_29 ipv4 | Enter the trunk mode for rsvp. |
| (config-trunk)#primary fast-reroute protection one-to-one | Configure primary fast-reroute protection facility for a trunk. |
| (config-trunk)#primary fast-reroute node-protection | Configure primary fast-reroute node protection for the trunk |
| (config-trunk)#primary path p41 | Configure trunk to_29 to use the defined path. |
| (config-trunk)#to 29.29.29.29 | Specify the IPv4 egress (destination point) for the LSP. |
| (config-trunk)#exit | Exit from trunk mode. |
| (config)#ip vrf vrf1 | Configure VRF instance |
| (config-vrf)#rd 100:1 | Configure Router Distinguisher value |
| (config-vrf)#route-target both 100:1 | Configure route-target as both |
| (config-vrf)#exit | Exit the path mode. |
| (config)#interface xe23 | Enter to the interface mode |
| (config-if)#ip vrf forwarding vrf1 | Bind the VRF instance to the interface |
| (config-if)#ip address 23.23.23.29/24 | Configure IP address |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 100 | Configure BGP router instance |
| (config-router)#neighbor 29.29.29.29 remote-as 100 | Configure neighbor with remote-as |
| (config-router)#neighbor 29.29.29.29 update-source 30.30.30.30 | Configure update source as loopback address |
| (config-router)#address-family vpnv4 unicast | Configure VPNv4 address family |
| (config-router-af)#neighbor 29.29.29.29 activate | Activate the VPN neighbor |
| (config-router-af)#exit-address-family | Exit the VPN address family |
| (config-router)#address-family ipv4 vrf vrf1 | Configure VRF address family |
| (config-router-af)redistribute connected | Redistribute connected route |
| (config-router-af)exit-address-family | Exit VRF address family |
| (config-router)#exit | Exit router mode |
| (config)#mpls l2-circuit vlan10 10 29.29.29.29 | Configure Virtual circuit. |
| (config-pseudowire)#exit | Exit pseudowire config mode. |
| (config)#service-template st1 | Template configuration |

| | |
|---|--|
| (config-svc)#match outer-vlan 10 | Match criteria under template configuration |
| (config-svc)#exit | Exit service template mode |
| (config)#service-template st2 | Template configuration |
| (config-svc)#match outer-vlan 30 | Match criteria under template configuration |
| (config-svc)#exit | Exit service template mode |
| (config)#interface xe24 | Enter interface configuration mode |
| (config-if)#switchport | Configure interface as switch port |
| (config-if)#mpls-l2-circuit vlan10 service-template st1 | Bind the interface to the VC with service template |
| (config-if)#exit | Exit interface configuration mode |
| (config)#mpls vpls vpls30 30 | Configure VPLS instance |
| (config-vpls)#signaling ldp | Configure VPLS signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Configure VPLS type as VLAN encapsulation |
| (config-vpls-sig)#vpls-peer 29.29.29.29 | Configure VPLS peer |
| (config-vpls-sig)#exit-signaling | Exit VPLS configuration mode |
| (config)#interface xe25 | Enter interface configuration mode |
| (config-if)#switchport | Configure interface as switch port |
| (config-if)#mpls-vpls vpls30 service-template st2 | Bind the VPLS instance to the interface |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config-if)#commit | Commit the transaction. |

Validation

```

RTR-30#show rsvp session
Ingress RSVP:
To          From          State          Pri Rt  Style Labelin Labelout
LSPName          Uptime  Est.time  D
SType
29.29.29.29      30.30.30.30      Up          Yes 1 1 SE    -      24322  to_29-
Primary          00:07:53 0s 118ms  D
EFAULT
29.29.29.29      69.69.69.30      Up          No  1 1 SE    -      24322  to_29-
Detour          00:07:53 0s 4ms  DEF
AULT
Total 2 displayed, Up 2, Down 0.

Egress RSVP:
To          From          State          Pri Rt  Style Labelin Labelout
LSPName          Uptime  Est.time  D
SType
30.30.30.30      29.29.29.29      Up          Yes 1 1 SE    24960  -      to_30-
Primary          00:07:57 N/A    ELSP
_CON
30.30.30.30      25.25.25.42      Up          Yes 1 1 SE    24961  -      to_30-
Detour          00:07:57 N/A    ELSP
_CON

```

Total 2 displayed, Up 2, Down 0.

RTR-30#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|----------------|--------|-------------|-----|----------|-----------|------|
| Intf | Nexthop | | | | | | |
| R(t)> | 29.29.29.29/32 | 1 | 5001 | Yes | LSP_ | | |
| DEFAULT | 24322 | xe2 | 41.41.41.31 | | | | |
| R(t)> | 29.29.29.29/32 | 2 | 5001 | No | LSP_ | | |
| DEFAULT | 24322 | xe1 | 69.69.69.42 | | | | |

RTR-30#

RTR-30#show mpls vrf-table

Output for IPv4 VRF table with id: 2

Primary FTN entry with FEC: 43.43.43.0/24, id: 1, row status: Active
 Owner: BGP, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 5001, Protected LSP id: 0, QoS Resource id: 0, Description: N/A
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
 Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 6, owner: BGP, out intf: xe1, out label: 25602
 Nexthop addr: 29.29.29.29 cross connect ix: 7, op code: Push and Lookup

Link 41.41.41.0/24 Goes down. Interface xe41 on router 30 is administratively disabled with the "shutdown command".

RTR-30#

RTR-30#configure terminal

Enter configuration commands, one per line. End with CNTL/Z.

RTR-30(config)#int xe41
 RTR-30(config-if)#shutdown
 RTR-30(config-if)#

RTR-30#show rsvp session

Ingress RSVP:

| To | From | State | Pri | Rt | Style | Labelin | Labelout |
|-------------|-------------|-----------------|-----|-----|-------|---------|----------|
| LSPName | | Uptime | D | | | | |
| SType | | Est.time | | | | | |
| 29.29.29.29 | 30.30.30.30 | Using Backup | Yes | 0 0 | SE | - | to_29- |
| Primary | N/A | DEFAULT | | | | | |
| 29.29.29.29 | 30.30.30.30 | Dn | Yes | 0 0 | SE | - | to_29- |
| Primary | N/A | DEFAULT | | | | | |
| 29.29.29.29 | 69.69.69.30 | Up | No | 1 1 | SE | 24322 | to_29- |
| Detour | | 00:10:53 0s 4ms | DEF | | | | |

AULT

Total 3 displayed, Up 1, Down 2.

Egress RSVP:

| To | From | State | Pri | Rt | Style | Labelin | Labelout |
|-------------|-------------|--------------|------|-----|-------|---------|----------|
| LSPName | | Uptime | D | | | | |
| SType | | Est.time | | | | | |
| 30.30.30.30 | 29.29.29.29 | Up | Yes | 1 1 | SE | 24960 | to_30- |
| Primary | | 00:10:57 N/A | ELSP | | | | |

_CON

Total 1 displayed, Up 1, Down 0.

RTR-30#show mpls vc-table

| VC-ID | Vlan-ID | Inner-Vlan-ID | Access-Intf | Network-Intf | Out Label | Tunnel- |
|-------|---------|---------------|-------------|--------------|-----------|-------------------|
| Label | Nexthop | Status | | | | |
| 10 | N/A | N/A | xe4 | xe1 | 24321 | 24322 29.29.29.29 |
| | | Active | | | | |

RTR-30#

RTR-30#show mpls vpls mesh

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW- |
|---------|--------------|--------------|----------|--------------|-----------|---------|-----|
| INDEX | SIG-Protocol | Status | | | | | |
| 30 | 29.29.29.29 | 24322 | 24320 | xe41 | 24320 | 2/Up | 2 L |
| DP | Active | | | | | | |

Link 41.41.41.0/24 is reestablished. Interface xe41 is administratively re-enabled.

RTR-30#configure terminal

Enter configuration commands, one per line. End with CNTL/Z.

RTR-30(config)#int xe41

RTR-30(config-if)#no shutdown

RTR-30(config-if)#

RTR-30#

RTR-30#show rsvp session

Ingress RSVP:

| To | From | State | Uptime | Est.time | Pri | Rt | Style | Labelin | Labelout |
|-------------|-------------|-------|----------|----------|-----|-----|--------|---------|--------------|
| LSPName | | | | | D | | | | |
| SType | | | | | | | | | |
| 29.29.29.29 | 30.30.30.30 | Up | 00:00:01 | 0s 8ms | DEF | Yes | 1 1 SE | - | 24322 to_29- |
| Primary | | | | | | | | | |
| AULT | | | | | | | | | |
| 29.29.29.29 | 69.69.69.30 | Up | 00:00:01 | 0s 8ms | DEF | No | 1 1 SE | - | 24322 to_29- |
| Detour | | | | | | | | | |
| AULT | | | | | | | | | |

Total 2 displayed, Up 2, Down 0.

Egress RSVP:

| To | From | State | Uptime | Est.time | Pri | Rt | Style | Labelin | Labelout |
|-------------|-------------|-------|----------|----------|------|-----|--------|---------|----------|
| LSPName | | | | | D | | | | |
| SType | | | | | | | | | |
| 30.30.30.30 | 29.29.29.29 | Up | 00:13:22 | N/A | ELSP | Yes | 1 1 SE | 24960 | - to_30- |
| Primary | | | | | | | | | |
| _CON | | | | | | | | | |
| 30.30.30.30 | 25.25.25.42 | Up | 00:00:08 | N/A | ELSP | Yes | 1 1 SE | 24961 | - to_30- |
| Detour | | | | | | | | | |
| _CON | | | | | | | | | |

Total 2 displayed, Up 2, Down 0.

RTR-30#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|----------------|--------|-------------|-----|----------|-----------|------|
| Intf | Nexthop | | | | | | |
| R(t)> | 29.29.29.29/32 | 1 | 5001 | Yes | LSP_ | | |
| DEFAULT | 24322 | xe41 | 41.41.41.31 | | | | |
| R(t)> | 29.29.29.29/32 | 2 | 5001 | No | LSP_ | | |
| DEFAULT | 24322 | xe54/1 | 69.69.69.42 | | | | |



Note: The primary LSP, which is in using backup state shall continue to use backup path in case where secondary is provisioned after the LSP state is changed to switch to backup.

RSVP Detour Over Ring Topology

Overview


In OcNOS, this feature allows the detour formation in the ring topology to enhance the routing experience. The detour formation is a local protection mechanism to reroute the data traffic when a failure or congestion occurs in the primary Label Switched Path (LSP). In Multiprotocol Label Switching (MPLS), the primary LSP is the default path through which the data travels from the source to the destination node.

Feature Characteristics

This feature allows detour to take the upstream path of protected LSP, allowing a detour based protection in a ring topology. The upstream path of the protected LSP is the section of the network that precedes the PLR node in the network. This feature works for both path and sender-template method of detour formation. For the inter-op solutions that do not support the sender-template method, use the path method of detour formation.

In the below diagram, the data traffic path highlighted in green dots is the primary LSP. The link shown with the red cross is locally protected at the Point of Local Repair (PLR) node. A PLR node is a network device that reacts and takes action when a link fails. For continued data traffic flow, detour occurs through the red dotted line. Detour in MPLS is an alternate path used when the primary LSP encounters disruption or congestion.

Figure 33. RSVP-TE FRR failover ring topology Feature Characteristics



Benefits

This feature helps detour the data traffic when there is a link or node failure, keeping the data traffic loss to a minimum (less than 50ms when BFD negotiated for fastest detection).

Prerequisite

Before the detour configuration in a ring topology, configure the RSVP tunnel with fast reroute protection of the one-to-one method.

For more information, refer to the [Fast Reroute Configuration \(One-to-one Method\) \(page 758\)](#) section of the [RSVP-TE Configuration \(page 643\)](#) chapter in the OcNOS Multi-Protocol Label Switching Guide.

Configuration


This section shows the configuration procedure to create a detour in the ring topology.

Topology

Configure the primary LSP in the below ring topology from the head end to the tail end.

For example, consider PE1 as the head end and PE2 as the tail end, and the primary LSP is via R1, R2, and R3. In this case, first configure the [Fast Reroute Configuration \(One-to-one Method\) \(page 758\)](#) on the PE1 and PE2 and then configure the [detour-allow-primary-upstream-path \(page 788\)](#) command in all the nodes. For example, if the link between R3 and PE2 is down, the detour follows via primary LSP to reach PE2.

Figure 34. RSVP-TE FRR failover ring topology - 1:1 Detour



PE1 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

| | |
|--|--|
| PE1#configure terminal | Enter configure mode. |
| PE1 (config)#interface xe3 | Enter interface mode xe3. |
| PE1 (config-if)#ip address 61.61.61.3/24 | Configure IPv4 address 61.61.61.3.24. |
| PE1 (config-if)#label-switching | Configure label switching on xe3. |
| PE1 (config-if)#enable-rsvp | Enable RSVP on xe3. |
| PE1 (config-if)#exit | Exit interface mode. |
| PE1 (config)#interface xe12 | Enter interface mode xe12. |
| PE1 (config-if)#ip address 58.58.58.2/24 | Configure IPv4 address 58.58.58.2/24. |
| PE1 (config-if)#label-switching | Configure label switching on xe12. |
| PE1 (config-if)#enable-rsvp | Enable RSVP on xe12. |
| PE1 (config-if)#exit | Exit interface mode. |
| PE1 (config)#interface lo | Enter loopback interface mode. |
| PE1 (config-if)#ip address 26.26.26.26/32 secondary | Configure IPv4 address 26.26.26.26/32. |
| PE1 (config-if)#exit | Exit interface mode. |
| PE1 (config)#router ospf 100 | Enter OSPF router mode. |
| PE1 (config-router)#ospf router-id 26.26.26.26 | Assign router ID 26.26.26.26 for OSPF. |
| PE1 (config-router)#network 26.26.26.26/32 area 0.0.0.0 | Define network 26.26.26.26/32 under router OSPF. |
| PE1 (config-router)#network 58.58.58.0/24 area 0.0.0.0 | Define network 58.58.58.0/24 under router OSPF. |
| PE1 (config-router)#network 61.61.61.0/24 area 0.0.0.0 | Define network 61.61.61.0/24 under router OSPF. |
| PE1 (config-router)#exit | Exit router OSPF mode. |
| PE1 (config)#commit | Commit the transaction. |
| PE1 (config)#exit | Exit the configure mode. |

PE1 - RSVP Configurations

This section shows:

- The configuration of detour to take the upstream path of protected LSP.
- The configuration of the primary LSP and attaching it to the RSVP trunk.
- The configuration of the FRR.

| | |
|--|--|
| PE1#configure terminal | Enter configure mode. |
| PE1(config)#router rsvp | Enable RSVP globally. |
| PE1(config-router)#detour-allow-primary-upstream-path | Configure this CLI to allow detour to take primary upstream path. |
| PE1(config-router)#exit | Exit router RSVP mode. |
| PE1(config)#rsvp-path PE1-PE2-01 mpls | Configure RSVP path PE1-PE2-01 and enter path mode. |
| PE1(config-path)#61.61.61.2 strict | Configure this explicit route path as a strict hop. |
| PE1(config-path)#23.23.23.3 strict | Configure this explicit route path as a strict hop. |
| PE1(config-path)#41.41.41.3 strict | Configure this explicit route path as a strict hop. |
| PE1(config-path)#56.56.56.3 strict | Configure this explicit route path as a strict hop. |
| PE1(config-path)#rsvp-trunk TR-PE1-PE2-MP-01 ipv4 | Create an RSVP trunk TR-PE1-PE2-MP-01 and enter the trunk mode. |
| PE1(config-trunk)#primary fast-reroute protection one-to-one | Configure primary fast reroute protection. |
| PE1(config-trunk)#primary fast-reroute node-protection | Configure node protection. |
| PE1(config-trunk)#primary path PE1-PE2-01 | Configure trunk PE1-PE2-01 to use as the primary LSP. |
| PE1(config-trunk)#from 26.26.26.26 | Assign the source loopback address 26.26.26.26 to the RSVP trunk. |
| PE1(config-trunk)#to 22.22.22.22 | Assign the destination loopback address 22.22.22.22 to the RSVP trunk. |
| PE1(config-trunk)#exit | Exit router RSVP trunk mode. |
| PE1(config)#commit | Commit the transaction. |
| PE1(config)#exit | Exit the configure mode. |

R1 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

| | |
|--|---------------------------------------|
| R1#configure terminal | Enter configure mode. |
| R1(config)#interface xe3 | Enter interface mode xe3. |
| R1(config-if)#ip address 61.61.61.2/24 | Configure IPv4 address 61.61.61.2/24. |
| R1(config-if)#label-switching | Configure label switching on xe3. |
| R1(config-if)#enable-rsvp | Enable RSVP on interface xe3. |
| R1(config-if)#exit | Exit interface mode. |

| | |
|---|---|
| R1 (config)#interface xe47 | Enter interface mode xe47. |
| R1 (config-if)#ip address 23.23.23.2/24 | Configure IPv4 address 23.23.23.2/24. |
| R1 (config-if)#label-switching | Configure label switching on xe47. |
| R1 (config-if)#enable-rsvp | Enable RSVP on interface xe47. |
| R1 (config-if)#exit | Exit interface mode. |
| R1 (config)#interface lo | Enter loopback interface mode. |
| R1 (config-if)#ip address 24.24.24.24/32 secondary | Configure IPv4 address 24.24.24.24/32. |
| R1 (config-if)#exit | Exit interface mode. |
| R1 (config)#router ospf 100 | Enter OSPF router mode. |
| R1 (config-router)#ospf router-id 24.24.24.24 | Assign router-id for OSPF. |
| R1 (config-router)#network 23.23.23.0/24 area 0.0.0.0 | Define network 23.23.23.0/24 under router OSPF. |
| R1 (config-router)#network 24.24.24.24/32 area 0.0.0.0 | Define network 24.24.24.24/32 under router OSPF. |
| R1 (config-router)#network 61.61.61.0/24 area 0.0.0.0 | Define network 61.61.61.0/24 under router OSPF. |
| R1 (config-router)#exit | Exit router OSPF mode. |
| R1 (config)#commit | Commit the transaction. |
| R1 (config)#exit | Exit the configure mode. |

R1 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

| | |
|---|--|
| R1#configure terminal | Enter configure mode. |
| R1 (config)#router rsvp | Enable RSVP globally. |
| R1 (config-router)#detour-allow-primary- upstream-path | Configure this CLI to allow detour to take primary upstream path. |
| R1 (config-router)#exit | Exit router RSVP mode. |
| R1 (config)#commit | Commit the transaction. |
| R1 (config)#exit | Exit the configure mode. |

R2 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

| | |
|-----------------------|-----------------------|
| R2#configure terminal | Enter configure mode. |
|-----------------------|-----------------------|

| | |
|---|---|
| R2 (config)#interface xe1 | Enter interface mode xe1. |
| R2 (config-if)#ip address 41.41.41.2/24 | Configure IPv4 address 41.41.41.2/24. |
| R2 (config-if)#label-switching | Configure label switching on xe1. |
| R2 (config-if)#enable-rsvp | Enable RSVP on xe1. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#interface xe47 | Enter interface mode xe47. |
| R2 (config-if)#ip address 23.23.23.3/24 | Configure IPv4 address 23.23.23.3/24. |
| R2 (config-if)#label-switching | Configure label switching on xe47. |
| R2 (config-if)#enable-rsvp | Enable RSVP on xe47. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#interface lo | Enter loopback interface mode. |
| R2 (config-if)#ip address 88.88.88.88/32 secondary | Configure IPv4 address 88.88.88.88/32. |
| R2 (config-if)#exit | Exit interface mode. |
| R2 (config)#router ospf 100 | Enter OSPF router mode. |
| R2 (config-router)#ospf router-id 88.88.88.88 | Assign router-id 88.88.88.88 for OSPF. |
| R2 (config-router)#network 23.23.23.0/24 area 0.0.0.0 | Define network 23.23.23.0/24 under router OSPF. |
| R2 (config-router)#network 41.41.41.0/24 area 0.0.0.0 | Define network 41.41.41.0/24 under router OSPF. |
| R2 (config-router)#network 88.88.88.88/32 area 0.0.0.0 | Define network 88.88.88.88/32 under router OSPF. |
| R2 (config-router)#exit | Exit router OSPF mode. |
| R2 (config)#commit | Commit the transaction. |
| R2 (config)#exit | Exit the configure mode. |

R2 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

| | |
|---|--|
| R2#configure terminal | Enter configure mode. |
| R2 (config)#router rsvp | Enable RSVP globally. |
| R2 (config-router)#detour-allow-primary- upstream-path | Configure this CLI to allow detour to take primary upstream path. |
| R2 (config-router)#exit | Exit router RSVP mode. |
| R2 (config)#commit | Commit the transaction. |
| R2 (config)#exit | Exit the configure mode. |

R3 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

| | |
|--|---|
| R3#configure terminal | Enter configure mode. |
| R3(config)#interface xe1 | Enter interface mode xe1. |
| R3(config-if)#ip address 41.41.41.3/24 | Configure IPv4 address 41.41.41.3/24. |
| R3(config-if)#label-switching | Configure label switching on xe1. |
| R3(config-if)#enable-rsvp | Enable RSVP on xe1. |
| R3(config-if)#exit | Exit interface mode. |
| R3(config)#interface xe5 | Enter interface mode xe5. |
| R3(config-if)#ip address 56.56.56.2/24 | Configure IPv4 address 56.56.56.2/24. |
| R3(config-if)#label-switching | Configure label switching on xe5. |
| R3(config-if)#enable-rsvp | Enable RSVP on xe5. |
| R3(config-if)#exit | Exit interface mode. |
| R3(config)#interface lo | Enter loopback interface mode. |
| R3(config-if)#ip address 99.99.99.99/32 secondary | Configure IPv4 address 99.99.99.99/32. |
| R3(config-if)#exit | Exit interface mode. |
| R3(config)#router ospf 100 | Enter OSPF router mode. |
| R3(config-router)#ospf router-id 99.99.99.99 | Assign router-id for OSPF. |
| R3(config-router)#network 41.41.41.0/24 area 0.0.0.0 | Define network 41.41.41.0/24 under router OSPF. |
| R3(config-router)#network 56.56.56.0/24 area 0.0.0.0 | Define network 56.56.56.0/24 under router OSPF. |
| R3(config-router)#network 99.99.99.99/32 area 0.0.0.0 | Define network 99.99.99.99/32 under router OSPF. |
| R3(config-router)#exit | Exit router OSPF mode. |
| R3(config)#commit | Commit the transaction. |
| R3(config)#exit | Exit the configure mode. |

R3 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

| | |
|---|--|
| R3#configure terminal | Enter configure mode. |
| R3(config)#router rsvp | Enable RSVP globally. |
| R3(config-router)#detour-allow-primary- | Configure this CLI to allow detour to take primary |

| | |
|--------------------------|--------------------------|
| upstream-path | upstream path. |
| R3 (config-router) #exit | Exit router RSVP mode. |
| R3 (config) #commit | Commit the transaction. |
| R3 (config) #exit | Exit the configure mode. |

R5 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

| | |
|--|---|
| R5#configure terminal | Enter configure mode. |
| R5 (config) #interface xe1 | Enter interface mode 58.58.58.3/24. |
| R5 (config-if) #ip address 58.58.58.3/24 | Configure IPv4 address. |
| R5 (config-if) #label-switching | Configure label switching on xe1. |
| R5 (config-if) #enable-rsvp | Enable RSVP on xe1. |
| R5 (config-if) #exit | Exit interface mode. |
| R5 (config) #interface xe13 | Enter interface mode xe13. |
| R5 (config-if) #ip address 54.54.54.4/24 | Configure IPv4 address 54.54.54.4/24. |
| R5 (config-if) #label-switching | Configure label switching on xe13. |
| R5 (config-if) #enable-rsvp | Enable RSVP on xe13. |
| R5 (config-if) #exit | Exit interface mode. |
| R5 (config) #interface lo | Enter loopback interface mode. |
| R5 (config-if) #ip address 17.17.17.17/32 secondary | Configure IPv4 address 17.17.17.17/32. |
| R5 (config-if) #exit | Exit interface mode. |
| R5 (config) #router ospf 100 | Enter OSPF router mode. |
| R5 (config-router) #ospf router-id 17.17.17.17 | Assign router-id for OSPF. |
| R5 (config-router) #network 17.17.17.17/32 area 0.0.0.0 | Define network 17.17.17.17/32 under router OSPF. |
| R5 (config-router) #network 54.54.54.0/24 area 0.0.0.0 | Define network 54.54.54.0/24 under router OSPF. |
| R5 (config-router) #network 58.58.58.0/24 area 0.0.0.0 | Define network 58.58.58.0/24 under router OSPF. |
| R5 (config-router) #exit | Exit router OSPF mode. |
| R5 (config) #commit | Commit the transaction. |
| R5 (config) #exit | Exit the configure mode. |

R5 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

| | |
|--|--|
| R5#configure terminal | Enter configure mode. |
| R5(config)#router rsvp | Enable RSVP globally. |
| R5(config-router)#detour-allow-primary-upstream-path | Configure this CLI to allow detour to take primary upstream path |
| R5(config-router)#exit | Exit router RSVP mode |
| R5(config)#commit | Commit the transaction. |
| R5(config)#exit | Exit the configure mode. |

R4 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

| | |
|--|--|
| R4#configure terminal | Enter configure mode. |
| R4(config)#interface xe13 | Enter interface mode xe13. |
| R4(config-if)#ip address 54.54.54.3/24 | Configure IPv4 address 54.54.54.3/24. |
| R4(config-if)#label-switching | Configure label switching on xe13. |
| R4(config-if)#enable-rsvp | Enable RSVP on interface xe13. |
| R4(config-if)#exit | Exit interface mode. |
| R4(config)#interface xe8 | Enter interface mode xe8. |
| R4(config-if)#ip address 62.62.62.3/24 | Configure IPv4 address 62.62.62.3/24. |
| R4(config-if)#label-switching | Configure label switching on xe8. |
| R4(config-if)#enable-rsvp | Enable RSVP on xe8. |
| R4(config-if)#exit | Exit interface mode. |
| R4(config)#interface lo | Enter loopback interface mode. |
| R4(config-if)#ip address 48.48.48.48/32 secondary | Configure IPv4 address 48.48.48.48/32. |
| R4(config-if)#exit | Exit interface mode. |
| R4(config)#router ospf 100 | Enter OSPF router mode. |
| R4(config-router)#ospf router-id 48.48.48.48 | Assign router-id for OSPF. |
| R4(config-router)#network 48.48.48.48/32 area 0.0.0.0 | Define network 48.48.48.48/32 under router OSPF. |
| R4(config-router)#network 54.54.54.0/24 area 0.0.0.0 | Define network 54.54.54.0/24 under router OSPF. |
| R4(config-router)#network 62.62.62.0/24 | Define network 62.62.62.0/24 under router OSPF. |

| | |
|------------------------|--------------------------|
| area 0.0.0.0 | |
| R4(config-router)#exit | Exit router OSPF mode. |
| R4(config)#commit | Commit the transaction. |
| R4(config)#exit | Exit the configure mode. |

R4 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

| | |
|--|---|
| R4#configure terminal | Enter configure mode. |
| R4(config)#router rsvp | Enable RSVP globally. |
| R4(config-router)#detour-allow-primary-upstream-path | Configure this CLI to allow detour to take primary upstream path. |
| R4(config-router)#exit | Exit router RSVP mode. |
| R4(config)#commit | Commit the transaction. |
| R4(config)#exit | Exit the configure mode. |

PE2 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

| | |
|---|--|
| PE2#configure terminal | Enter configure mode. |
| PE2(config)#interface xe5 | Enter interface mode xe5. |
| PE2(config-if)#ip address 56.56.56.3/24 | Configure IPv4 address 56.56.56.3/24. |
| PE2(config-if)#label-switching | Configure label switching on xe5. |
| PE2(config-if)#enable-rsvp | Enable RSVP on xe5. |
| PE2(config-if)#exit | Exit interface mode. |
| PE2(config)#interface ge8 | Enter interface mode ge8. |
| PE2(config-if)#ip address 62.62.62.2/24 | Configure IPv4 address 62.62.62.2/24. |
| PE2(config-if)#label-switching | Configure label switching on ge8. |
| PE2(config-if)#enable-rsvp | Enable RSVP on ge8. |
| PE2(config-if)#exit | Exit interface mode. |
| PE2(config)#interface lo | Enter loopback interface mode. |
| PE2(config-if)#ip address 22.22.22.22/32 secondary | Configure IPv4 address 22.22.22.22/32. |
| PE2(config-if)#exit | Exit interface mode. |
| PE2(config)#router ospf 100 | Enter OSPF router mode. |
| PE2(config-router)#ospf router-id | Assign router-id for OSPF. |

| | |
|---|--|
| 22.22.22.22 | |
| PE2 (config-router) #network 22.22.22.22/32 area 0.0.0.0 | Define network 22.22.22.22/32 under router OSPF. |
| PE2 (config-router) #network 56.56.56.0/24 area 0.0.0.0 | Define network 56.56.56.0/24 under router OSPF. |
| PE2 (config-router) #network 62.62.62.0/24 area 0.0.0.0 | Define network 62.62.62.0/24 under router OSPF. |
| PE2 (config-router) #exit | Exit router OSPF mode. |
| PE2 (config) #commit | Commit the transaction. |
| PE2 (config) #exit | Exit the configure mode. |

PE2 - RSVP Configurations

This section shows:

- The configuration of detour to take the upstream path of protected LSP.
- The configuration of the primary LSP and attaching it to the RSVP trunk.
- The configuration of the FRR.

| | |
|--|---|
| PE2#configure terminal | Enter configure mode. |
| PE2 (config) #router rsvp | Enable RSVP globally. |
| PE2 (config-router) #detour-allow-primary-upstream-path | Configure this CLI to allow detour to take primary upstream path. |
| PE2 (config-router) #exit | Exit router RSVP mode. |
| PE2 (config) #rsvp-path PE2-PE1-01 mpls | Configure RSVP path PE2-PE1-01 and enter path mode. |
| PE2 (config-path) #56.56.56.2 strict | Configure this explicit route path as a strict hop. |
| PE2 (config-path) #41.41.41.2 strict | Configure this explicit route path as a strict hop. |
| PE2 (config-path) #23.23.23.2 strict | Configure this explicit route path as a strict hop. |
| PE2 (config-path) #61.61.61.3 strict | Configure this explicit route path as a strict hop. |
| PE2 (config-router) #exit | Exit path mode. |
| PE2 (config-path) #rsvp-trunk TR-PE2-PE1-MP-01 ipv4 | Create an RSVP trunk TR-PE2-PE1-MP-01 and enter the Trunk mode. |
| PE2 (config-trunk) #primary fast-reroute protection one-to-one | Configure primary fast-reroute protection. |
| PE2 (config-trunk) #primary fast-reroute node-protection | Configure node protection. |
| PE2 (config-trunk) #primary path PE2-PE1-01 | Configure trunk PE2-PE1-01 to use as the primary LSP. |
| PE2 (config-trunk) #from 22.22.22.22 | Assign the source loopback address 22.22.22.22 |

| | |
|-----------------------------------|--|
| | to the RSVP trunk. |
| PE2 (config-trunk)#to 26.26.26.26 | Assign the destination loopback address 26.26.26.26 to the RSVP trunk. |
| PE2 (config-trunk)#exit | Exit router RSVP trunk mode. |
| PE2 (config)#commit | Commit the transaction. |
| PE2 (config)#exit | Exit the configure mode. |

Validation

PE1

Below is the validation output of RSVP LSPs from PE1 to PE2 via R1>R2>R3:

```
#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt    Style  Labelin  Labelout
22.22.22.22 26.26.26.26 5001      2205      PRI     TR-PE1-PE2-MP-01-
Primary     UP    02:12:32 1 1 SE    -       52480
22.22.22.22 58.58.58.2 5001      2205      DTR     TR-PE1-PE2-MP-01-
Detour      UP    00:34:04 1 2 SE    -       25600
Total 2 displayed, Up 2, Down 0.

Transit RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt    Style  Labelin  Labelout
22.22.22.22 61.61.61.2 5001      2205      PRI     TR-PE1-PE2-MP-01-
Detour      UP    00:33:19 1 2 SE    25602    25600
Total 1 displayed, Up 1, Down 0.

Egress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt    Style  Labelin  Labelout
26.26.26.26 22.22.22.22 5001      2205      PRI     TR-PE2-PE1-MP-01-
Primary     UP    02:12:27 1 1 SE    25601    -
26.26.26.26 62.62.62.2 5001      2205      PRI     TR-PE2-PE1-MP-01-
Detour      UP    02:09:08 1 1 SE    25600    -
Total 2 displayed, Up 2, Down 0.
```

Below is the validation output of RSVP ping and trace from PE1 to PE2:

```
#ping mpls rsvp egress 22.22.22.22 detail
Sending 5 MPLS Echos to 22.22.22.22, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

! seq_num = 1 56.56.56.3 0.91 ms
```

```

! seq_num = 2 56.56.56.3 0.54 ms
! seq_num = 3 56.56.56.3 0.48 ms
! seq_num = 4 56.56.56.3 0.47 ms
! seq_num = 5 56.56.56.3 0.50 ms

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.47/0.69/0.91
PE1#
#trace mpls rsvp egress 22.22.22.22 detail
Tracing MPLS Label Switched Path to 22.22.22.22, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

  0 61.61.61.3 [Labels: 52480]
R 1 61.61.61.2 [Labels: 25600] 0.71 ms
R 2 23.23.23.3 [Labels: 25600] 0.83 ms
R 3 41.41.41.3 [Labels: 25600] 0.88 ms
! 4 56.56.56.3 0.69 ms

```

Below are the outputs from transit nodes R1, R2, and R3 for primary LSP configured:

R1

```

#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime     Rt   Style  Labelin  Labelout
22.22.22.22 61.61.61.2 5001     2205     DTR    TR-PE1-PE2-MP-01-
Detour      UP    00:38:43 1 2 SE    -      25602
26.26.26.26 23.23.23.2 5001     2205     DTR    TR-PE2-PE1-MP-01-
Detour      UP    00:38:44 1 1 SE    -      25603
Total 2 displayed, Up 2, Down 0.

Transit RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime     Rt   Style  Labelin  Labelout
22.22.22.22 26.26.26.26 5001     2205     PRI    TR-PE1-PE2-MP-01-
Primary     UP    02:17:55 1 1 SE    52480  25600
22.22.22.22 23.23.23.3 5001     2205     PRI    TR-PE1-PE2-MP-01-
Detour      UP    00:37:58 1 2 SE    52482  25602
26.26.26.26 22.22.22.22 5001     2205     PRI    TR-PE2-PE1-MP-01-
Primary     UP    02:17:50 1 1 SE    52481  25601
Total 3 displayed, Up 3, Down 0.

```

R2

```

#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes

```

(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-------------|--------|--------|-------|-------------------|-------|
| 22.22.22.22 | 23.23.23.3 | 5001 | 2205 | DTR | TR-PE1-PE2-MP-01- | |
| Detour | UP 00:38:07 | 1 2 SE | - | 52482 | | |
| 26.26.26.26 | 41.41.41.2 | 5001 | 2205 | DTR | TR-PE2-PE1-MP-01- | |
| Detour | UP 00:39:00 | 1 2 SE | - | 25602 | | |

Total 2 displayed, Up 2, Down 0.

Transit RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-------------|--------|--------|-------|-------------------|-------|
| 22.22.22.22 | 26.26.26.26 | 5001 | 2205 | PRI | TR-PE1-PE2-MP-01- | |
| Primary | UP 02:18:05 | 1 1 SE | 25600 | 25600 | | |
| 22.22.22.22 | 41.41.41.3 | 5001 | 2205 | PRI | TR-PE1-PE2-MP-01- | |
| Detour | UP 00:37:28 | 1 2 SE | 25602 | 52482 | | |
| 26.26.26.26 | 22.22.22.22 | 5001 | 2205 | PRI | TR-PE2-PE1-MP-01- | |
| Primary | UP 02:18:00 | 1 1 SE | 25601 | 52481 | | |
| 26.26.26.26 | 23.23.23.2 | 5001 | 2205 | PRI | TR-PE2-PE1-MP-01- | |
| Detour | UP 00:38:53 | 1 2 SE | 25603 | 25602 | | |

Total 4 displayed, Up 4, Down 0.

R3

```
#show rsvp session
```

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
 State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
 * indicates the session is active with local repair at one or more nodes
 (P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-------------|--------|--------|-------|-------------------|-------|
| 22.22.22.22 | 41.41.41.3 | 5001 | 2205 | DTR | TR-PE1-PE2-MP-01- | |
| Detour | UP 00:37:31 | 1 1 SE | - | 25602 | | |
| 26.26.26.26 | 56.56.56.2 | 5001 | 2205 | DTR | TR-PE2-PE1-MP-01- | |
| Detour | UP 00:39:23 | 1 2 SE | - | 25602 | | |

Total 2 displayed, Up 2, Down 0.

Transit RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-------------|--------|--------|-------|-------------------|-------|
| 22.22.22.22 | 26.26.26.26 | 5001 | 2205 | PRI | TR-PE1-PE2-MP-01- | |
| Primary | UP 02:18:08 | 1 1 SE | 25600 | 25600 | | |
| 26.26.26.26 | 22.22.22.22 | 5001 | 2205 | PRI | TR-PE2-PE1-MP-01- | |
| Primary | UP 02:18:02 | 1 1 SE | 25601 | 25601 | | |
| 26.26.26.26 | 41.41.41.2 | 5001 | 2205 | PRI | TR-PE2-PE1-MP-01- | |
| Detour | UP 00:39:03 | 1 2 SE | 25602 | 25602 | | |

Total 3 displayed, Up 3, Down 0.

Below are the outputs from transit nodes R4 and R5 for Detour LSPs formation:

From R4

```
#show rsvp session
```

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
 State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
 * indicates the session is active with local repair at one or more nodes
 (P) indicates the secondary-priority session is acting as primary

Transit RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|------------|--------|--------|------|-------------------|-------|
| 22.22.22.22 | 58.58.58.2 | 5001 | 2205 | PRI | TR-PE1-PE2-MP-01- | |

```

Detour      UP      02:14:52  1 1 SE      25600      25601
26.26.26.26 62.62.62.2      5001      2205      PRI      TR-PE2-PE1-MP-01-
Detour      UP      00:39:49  1 1 SE      25601      25601
Total 2 displayed, Up 2, Down 0.

```

From R5

```

#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

```

Transit RSVP:
To      From      Tun-ID      LSP-ID      Type      LSPName      State
Uptime  Rt  Style  Labelin  Labelout
22.22.22.22 58.58.58.2 5001      2205      PRI      TR-PE1-PE2-MP-01-
Detour      UP      00:39:45  1 1 SE      25600      25600
26.26.26.26 62.62.62.2 5001      2205      PRI      TR-PE2-PE1-MP-01-
Detour      UP      02:14:48  1 1 SE      25601      25601
Total 2 displayed, Up 2, Down 0.

```

Now, shutting down one of the interfaces on Primary LSP path and check RSVP tunnel outputs on PE1 and PE2
Shutdown interface xe47 connected between R1 and R2:

| | |
|--------------------------|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface xe47 | Enter interface mode. |
| (config-router)#shutdown | Administratively bring the interface down. |
| (config-router)#exit | Exit router RSVP mode |

Below is the validation output of RSVP LSPs from PE1 to PE2 after admin shutting one of the interfaces on primary LSP path:

```

#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

```

Ingress RSVP:
To      From      Tun-ID      LSP-ID      Type      LSPName      State
Uptime  Rt  Style  Labelin  Labelout
22.22.22.22 26.26.26.26 5001      2205      PRI      TR-PE1-PE2-MP-01-
Primary  UP*  02:32:40  1 1 SE      -      52480
22.22.22.22 26.26.26.26 5001      2201      PRI      TR-PE1-PE2-MP-01-
Primary  DN   N/A      0 0 SE      -
22.22.22.22 58.58.58.2 5001      2205      DTR      TR-PE1-PE2-MP-01-
Detour      UP      00:54:12  1 2 SE      -      25600
Total 3 displayed, Up 2, Down 1.

Transit RSVP:
To      From      Tun-ID      LSP-ID      Type      LSPName      State
Uptime  Rt  Style  Labelin  Labelout
22.22.22.22 61.61.61.2 5001      2205      PRI      TR-PE1-PE2-MP-01-
Detour      UP      00:53:27  1 2 SE      25602      25600
Total 1 displayed, Up 1, Down 0.

```

Below is the validation output of RSVP ping and trace from PE1 to PE2 after shutting one of the interfaces on primary LSP path:

```
Egress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
26.26.26.26 62.62.62.2 5001      2205      PRI     TR-PE2-PE1-MP-01-
Detour      UP    02:29:16 1 1 SE    25600    -
Total 1 displayed, Up 1, Down 0.

#ping mpls rsvp egress 22.22.22.22 detail
Sending 5 MPLS Echos to 22.22.22.22, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

! seq_num = 1 62.62.62.2 0.69 ms
! seq_num = 2 62.62.62.2 0.54 ms
! seq_num = 3 62.62.62.2 0.56 ms
! seq_num = 4 62.62.62.2 0.49 ms
! seq_num = 5 62.62.62.2 0.51 ms

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.49/0.59/0.69
#trace mpls rsvp egress 22.22.22.22 detail
Tracing MPLS Label Switched Path to 22.22.22.22, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

  0 61.61.61.3 [Labels: 52480]
R 1 61.61.61.2 [Labels: 25602] 0.72 ms
R 2 61.61.61.3 [Labels: 25600] 0.67 ms
R 3 58.58.58.3 [Labels: 25600] 0.80 ms
R 4 54.54.54.3 [Labels: 25601] 0.80 ms
! 5 62.62.62.2 0.50 ms
```

Below is the validation output of RSVP LSPs from PE2 to PE1 after admin shutting one of the interfaces on primary LSP path:

```
#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
26.26.26.26 22.22.22.22 5001      2205      PRI     TR-PE2-PE1-MP-01-
Primary     UP*    02:36:19 1 1 SE    -        25601
```

```

26.26.26.26      22.22.22.22      5001      2201      PRI      TR-PE2-PE1-MP-01-
Primary         DN      N/A      0 0 SE      -
26.26.26.26      62.62.62.2      5001      2205      DTR      TR-PE2-PE1-MP-01-
Detour         UP      00:57:57  1 2 SE      -      25601
Total 3 displayed, Up 2, Down 1.

```

Transit RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|------------|----------|---------|----------|-------------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 26.26.26.26 | 56.56.56.2 | 5001 | 2205 | PRI | TR-PE2-PE1-MP-01- | |
| Detour | UP | 00:57:40 | 1 2 SE | 25602 | 25601 | |

Total 1 displayed, Up 1, Down 0.

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|------------|----------|---------|----------|-------------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 22.22.22.22 | 58.58.58.2 | 5001 | 2205 | PRI | TR-PE1-PE2-MP-01- | |
| Detour | UP | 02:33:00 | 1 1 SE | 25601 | - | |

Total 1 displayed, Up 1, Down 0.

Below is the validation output of RSVP ping and trace from PE2 to PE1 after shutting one of the interfaces on primary LSP path:

```

#ping mpls rsvp egress 26.26.26.26 detail
Sending 5 MPLS Echos to 26.26.26.26, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

! seq_num = 1 58.58.58.2 0.80 ms
! seq_num = 2 58.58.58.2 0.59 ms
! seq_num = 3 58.58.58.2 0.47 ms
! seq_num = 4 58.58.58.2 0.49 ms
! seq_num = 5 58.58.58.2 0.54 ms

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.47/0.63/0.80
#trace mpls rsvp egress 26.26.26.26 detail
Tracing MPLS Label Switched Path to 26.26.26.26, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

  0 56.56.56.3 [Labels: 25601]
R 1 56.56.56.2 [Labels: 25601] 1.01 ms
R 2 41.41.41.2 [Labels: 25602] 0.95 ms
R 3 41.41.41.3 [Labels: 25602] 0.62 ms
R 4 56.56.56.3 [Labels: 25601] 0.79 ms
R 5 62.62.62.3 [Labels: 25601] 0.67 ms
R 6 54.54.54.4 [Labels: 25600] 0.57 ms
! 7 58.58.58.2 0.50 ms

```

Implementation Examples

To implement detour based protection in a ring topology, use the command [detour-allow-primary-upstream-path \(page 788\)](#) that allows the detour formation to consider the upstream path of protected LSP. This is only applicable in ring topology.

RSVP Detour Over Ring Topology Command

detour-allow-primary-upstream-path

Use this command to ensure detour formation to consider the upstream path of protected LSPs. This is a deviation to RFC 4090 section 6.2 recommendation (<https://datatracker.ietf.org/doc/html/rfc4090>). This command is intended to be used in special cases where detour protection is required on ring topology if no alternate path is available.

Use the no parameter with this command to bypass the upstream path to the protected LSP when choosing a detour path.



Note: This command is intended to be used in ring topology if detour support is required at the cost of resource and link bandwidth. This command is not recommended to be configured otherwise.

Command Syntax

```
detour-allow-primary-upstream-path
no detour-allow-primary-upstream-path
```

Parameters

None

Default

By default, detour formation excludes the protected LSP upstream path as per RFC 4090 section 6.2 recommendations.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#detour-allow-primary-upstream-path
(config-router)#commit
(config-router)#no detour-allow-primary-upstream-path
(config-router)#commit
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

| Acronym | Description |
|---------|--------------------------|
| FRR | Fast Reroute |
| LSP | Label Switched Path |
| OSPF | Open Shortest Path First |
| PLR | Point of Local Repair |

Glossary

The following provides definitions for key terms used throughout this document:

| | |
|---------------------------------------|--|
| Detour formation in the ring topology | The detour formation in the ring topology is a mechanism to reroute the data traffic over the backup path when a failure or congestion occurs in the primary Label Switched Path (LSP). |
| PLR node | A PLR node is a network device that reacts and takes action when a link fails. |
| Primary LSP | The primary LSP is the default path of the forwarding data packets from the source device to the destination device. |
| Protected LSP | A protected LSP is a primary LSP with a backup path in an MPLS network. When there is an issue or a failure in the primary LSP, the traffic is rerouted through the backup path, protecting the primary LSP. |
| RSVP Tunnel | RSVP tunnels are logical paths through which data traffic traverses in an IP network. |
| Upstream path of the protected LSP | The upstream path of the protected LSP is the section of the network that precedes the PLR node in the network. |

Auto-Bandwidth with RSVP-TE

Overview

Automatic bandwidth allows to dynamically adjust bandwidth reservation based on the measured traffic. RSVP automatic bandwidth monitors the traffic rate on a Label Switched Path (LSP) and resizes the bandwidth to align it closely with the traffic in the tunnel. RSVP automatic bandwidth is configured on individual LSPs at every headend router.

Auto bandwidth can be added to an operational LSP at any time, but no bandwidth change occurs until a future trigger event or auto bandwidth profile configured with initial bandwidth or minimum bandwidth. Auto bandwidth may also be removed from an operational LSP at any time and this would re-signal the LSP with no bandwidth reservation.

Feature Characteristics

The characteristics of the RSVP auto-bandwidth are:

- RSVP-TE auto bandwidth provides the means to automatically adjust the bandwidth allocation for traffic engineering tunnels based on their measured traffic load.
- This feature samples the average output rate for each tunnel marked for automatic bandwidth adjustment. For each marked tunnel, this feature periodically adjusts the tunnel's allocated bandwidth to the largest eligible sample for the tunnel since the last adjustment.
- The frequency with which tunnel bandwidth is adjusted and the allowable range of adjustments should be configurable on a per-auto-bandwidth profile basis.
- In addition, the sampling interval and the interval over which to average tunnel traffic to obtain the average output rate is user-configurable on a per-auto-bandwidth profile basis.

**Notes:**

- Convergence on redundancy may require bidirectional traffic or MAC aging.
- The feature relies on stat_id allocation to tunnel entities, and there is a limit on the maximum number of stat_ids (which varies based on the chip variant). If a tunnel is not associated with a stat entity, traffic rate samples cannot be fetched for those tunnels.
- RSVP Graceful Restart is not supported for automatic bandwidth. When a GR is performed, RSVP will not store the current bandwidth for the reservation. It will use either the initial bandwidth (if configured), the minimum bandwidth, or the highest bandwidth of the on-boot sample (if auto-bandwidth-on-boot is configured).
- The auto bandwidth feature relies on the hardware's ability to collect tunnel traffic counters. In Qumran1 devices, the "hardware-profile statistics tunnel-lif enable" command must be enabled, and the system must be reloaded for the change to take effect. Without the tunnel statistics profile, auto bandwidth will not process traffic rates and will be ineffective. Note that, only 2 statistics profiles shall be configured as this is the hardware limitation.
- Auto bandwidth and manual bandwidth configurations are mutually exclusive. Auto bandwidth allows for configuring an initial bandwidth, which will be used as the session's initial bandwidth when auto bandwidth is associated with a trunk. If the initial bandwidth is not configured, the minimum bandwidth will be used to initialize the session bandwidth.
- For tunnels with only one hop, the no PHP (default config) must be set for the rate to be computed correctly.

Benefits

In large MPLS transport networks in service provider settings with this capability:

- The network can react faster to sudden bursts of traffic in near real-time and not rely on manual intervention.
- Effective use of bandwidth resources by minimizing the over-subscription/padding of LSP bandwidth.
- Maximizes the usage of available bandwidth and optimizes the network effectively to use preferred, shorter latency, paths first.

Prerequisites

Define Interfaces and Loopback Addresses

Configure Layer 3 interfaces, like port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 100.1.1.1/32 secondary  
  ipv6 address ::1/128  
!  
interface xel  
  ip address 1.1.1.1/24  
!
```

Configure IGP for Dynamic Routing

Configure IGP for dynamic routing by following the steps mentioned. This setup includes enabling ISIS for dynamic routing and configuring OSPF for the network.

ISIS Configuration

1. Enable ISIS on all nodes: Ensure that ISIS is running across the network to facilitate dynamic routing.
2. Define ISIS Router Instances: Set up instances to match loopback IP addresses.
3. Add Network Segments to ISIS Areas: This ensures proper route distribution.
4. Set up Neighbor Relationships: Use loopback IP addresses to establish these relationships for efficient route advertisement and convergence.

```
!  
router isis 1  
  is-type level-2-only  
  metric-style wide  
  mpls traffic-eng router-id 100.1.1.1  
  mpls traffic-eng level-2  
  capability cspf  
  dynamic-hostname  
  fast-reroute ti-lfa level-2 proto ipv4  
  net 49.0000.0000.0001.00  
  passive-interface lo  
!  
interface xel  
  isis network point-to-point  
  ip router isis 1  
!
```

OSPF Configuration

1. Configure OSPF Router ID: Assign a unique router ID for OSPF operations.
2. Define OSPF Networks: Include the loopback IP and other network segments in the OSPF area for route distribution.

```
!  
router ospf 100  
  ospf router-id 100.1.1.1  
  network 100.1.1.1/32 area 0.0.0.0  
  network 1.1.1.1/24 area 0.0.0.0  
!
```

Configure RSVP for Efficient Network Operation

Enable Resource Reservation Protocol (RSVP) on all nodes to optimize traffic routing and quality of service. RSVP reserves network resources along specified paths to enhance network performance and reliability.

```
!  
router rsvp  
!  
interface xel  
  label-switching  
  enable-rsvp  
!
```

Configure the RSVP Primary Path and Trunk

Establish a trunk is required on edge routers participating in label-switching using defined path. Configuring the RSVP path is optional.

```
!
rsvp-path PE1-PE2-1 mpls
  1.1.1.2 strict
  1.1.2.1 strict
!
rsvp-trunk PE1-PE2 ipv4
  reoptimize
  primary fast-reroute protection facility
  primary fast-reroute node-protection
  primary path PE1-PE2-1
  from 100.1.1.1
  to 100.1.1.3
!
```

Configuration for RSVP Auto-Bandwidth


Configure various nodes within the topology to set up a RSVP-Auto bypass tunnels.

Topology

The sample topology includes Edge Nodes (PE1 and PE2) and core Nodes (P1).

Primary path is defined via PE1-P1-PE2.

Figure 35. RSVP-Auto Bypass Tunnel Setup



PE1

1. Create auto-bandwidth Profile.

```
(config)# rsvp-auto-bandwidth AUTO-BW
(config-auto-bandwidth)# commit
```

2. Set the Sample interval & adjust interval.

```
(config-auto-bandwidth)# sample-interval 1
(config-auto-bandwidth)# adjust-interval 5
```

3. Set the minimum & maximum bandwidth rate.

```
(config-auto-bandwidth)# minimum-bandwidth 200m
(config-auto-bandwidth)# maximum-bandwidth 500m
```

4. Set the overflow-threshold & underflow-threshold.

```
(config-auto-bandwidth)# overflow-threshold absolute 100m
(config-auto-bandwidth)# underflow-threshold absolute 50m
```

5. Set the overflow & underflow limit.

```
(config-auto-bandwidth)# overflow-limit 2
(config-auto-bandwidth)# underflow-limit 2
```

6. Set the maximum number of consecutive times the average bandwidth can exceed the maximum threshold bandwidth before the exceed action is applied.

```
(config-auto-bandwidth)# maximum-bandwidth-exceed-limit 2
(config-auto-bandwidth)#maximum-bandwidth-exceed-action teardown
(config-auto-bandwidth)#commit
```

Running Configuration on PE1 Router

```
#show running-config rsvp
!
router rsvp
!
!
interface xe1
  enable-rsvp
!
interface xe2
  enable-rsvp
!
!
!
!
!
!
rsvp-auto-bandwidth AUTO-BW
  sample-interval 1
  adjust-interval 5
  minimum-bandwidth 200m
  maximum-bandwidth 500m
  overflow-threshold absolute 100m
  underflow-threshold absolute 50m
  overflow-limit 2
  underflow-limit 2
  maximum-bandwidth-exceed-limit 2
  maximum-bandwidth-exceed-action teardown
!
rsvp-trunk PE1-PE2 ipv4
  reoptimize
  primary fast-reroute protection facility
  primary fast-reroute node-protection
  primary path PE1-PE2-1
  auto-bandwidth AUTO-BW
  from 100.1.1.1
  to 100.1.1.3
!
```

Validation

Verify auto bandwidth adjustments information as below:

Send the sample rate with 200MBPS and verify the auto bandwidth adjustments as below:

```
#show rsvp trunk auto-bandwidth detail
Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2202, Egress: 100.1.1.3
```

```
-----
Sample Interval           : 1 minutes
Adjustment Interval      : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth : 0
Minimum Bandwidth        : 200m
Maximum Bandwidth        : 500m
Overflow Threshold Bandwidth : 100m
```

```

Underflow Threshold Bandwidth      : 50m
Overflow Threshold Activate Bandwidth : 0
Underflow Threshold Activate Bandwidth : 0
Overflow Limit                     : 2
Underflow Limit                    : 2
Max. Bandwidth Exceed Limit        : 2
-----
Max-BW-exceed-limit action         : teardown
Resignal-failure-action            : notify
Monitor Bandwidth                  : No
-----
Minimum Average Bandwidth          : 0
Maximum Average Bandwidth          : 202.8m
Total Overflow Count               : 0
Consecutive Overflow Count         : 0
Consecutive Eligible Overflow Count : 0
Total Underflow Count              : 0
Consecutive Underflow Count        : 0
Consecutive Eligible Underflow Count : 0
Max. Bandwidth Exceed Count        : 0
Teardown Count                     : 0
-----
Last Bandwidth                     : 0
Last Requested Bandwidth           : 0
Last Signaled Bandwidth            : 0
Current Bandwidth                   : 200m
Highest Bandwidth                   : 203m
-----
Time for Next Sample request       : 48 seconds
Time for Next Adjustment           : 2 minutes, 57 seconds
Time of Last Bandwidth Request     : N/A
Time of Last Bandwidth Signal      : N/A
Time of Last Adjustment            : N/A
Time of Highest Bandwidth Marked   : 2024 Jun 25 09:56:19
-----
Total Auto-Bandwidth Adjustments    : 0
Successful Adjustments              : 0
Failed Adjustments                  : 0
-----
Samples collected in the current adjustment cycle:
[Sample 1-5]      : 202.8m  202.6m

```

```
#show rsvp trunk auto-bandwidth
```

| Trunk-Name | Trunk | LSP | Last | Requested | Signaled | Current | Highest | Adjust- |
|------------------|-------|------|------|-----------|----------|---------|---------|---------|
| Time Last-Adjust | ID | ID | BW | BW | BW | BW | BW | Left |
| (sec) Time | | | | | | | | |
| PE1-PE2 | 5002 | 2202 | 0 | 0 | 0 | 200m | 203m | 176 |
| | | | | | | | | NA |

Overflow:

Current bandwidth is adjusted to 290.4MBPS. Then, send the sample rate which is more than overflow threshold i.e, 340.1MBPS , 377.2MBPS.
 As per the below output current bandwidth is more than overflow threshold bandwidth and consecutively two samples are received and it is more than the overflow limit.
 So current BW is adjusted to 377.2mbps after 2 consecutive samples collected as per the Maximum Average Bandwidth.

```
#show rsvp trunk auto-bandwidth detail
```

```
Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2202, Egress: 100.1.1.3
```

```
Sample Interval      : 1 minutes
```

```

Adjustment Interval           : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth     : 200m
Minimum Bandwidth             : 100m
Maximum Bandwidth             : 500m
Overflow Threshold Bandwidth  : 40m
Underflow Threshold Bandwidth : 30m
Overflow Threshold Activate Bandwidth : 0
Underflow Threshold Activate Bandwidth : 0
Overflow Limit                : 2
Underflow Limit               : 2
Max. Bandwidth Exceed Limit   : 2
-----
Max-BW-exceed-limit action    : teardown
Resignal-failure-action       : notify
Monitor Bandwidth             : No
-----
Minimum Average Bandwidth     : 0
Maximum Average Bandwidth     : 340.1m
Total Overflow Count          : 1
Consecutive Overflow Count    : 1
Consecutive Eligible Overflow Count : 1
Total Underflow Count         : 0
Consecutive Underflow Count   : 0
Consecutive Eligible Underflow Count : 0
Max. Bandwidth Exceed Count   : 0
Teardown Count                : 0
-----
Last Bandwidth                : 190.7m
Last Requested Bandwidth      : 290.4m
Last Signaled Bandwidth       : 290.4m
Current Bandwidth             : 290.4m
Highest Bandwidth             : 340.1m
-----
Time for Next Sample request   : 0 seconds
Time for Next Adjustment      : 0 seconds
Time of Last Bandwidth Request : 2024 Jun 25 10:28:32
Time of Last Bandwidth Signal  : 2024 Jun 25 10:28:32
Time of Last Adjustment       : 2024 Jun 25 10:28:32
Time of Highest Bandwidth Marked : 2024 Jun 25 10:29:35
-----
Total Auto-Bandwidth Adjustments : 4
Successful Adjustments           : 4
Failed Adjustments               : 0
-----
Samples collected in the current adjustment cycle:
  [Sample 1-5]      : 340.1m

#show rsvp trunk auto-bandwidth detail
Session: PE1-PE2_1-Primary, Tunnel-id: 5002, LSP-ID: 2203, Egress: 100.1.1.3
-----
Sample Interval           : 1 minutes
Adjustment Interval       : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth  : 200m
Minimum Bandwidth         : 100m
Maximum Bandwidth         : 500m
Overflow Threshold Bandwidth : 40m
Underflow Threshold Bandwidth : 30m
Overflow Threshold Activate Bandwidth : 0
Underflow Threshold Activate Bandwidth : 0
Overflow Limit            : 2
Underflow Limit           : 2
Max. Bandwidth Exceed Limit : 2
-----
Max-BW-exceed-limit action : teardown

```



```

Resignal-failure-action      : notify
Monitor Bandwidth           : No
-----
Minimum Average Bandwidth    : 0
Maximum Average Bandwidth    : 0
Total Overflow Count         : 0
Consecutive Overflow Count   : 0
Consecutive Eligible Overflow Count : 0
Total Underflow Count        : 0
Consecutive Underflow Count  : 0
Consecutive Eligible Underflow Count : 0
Max. Bandwidth Exceed Count  : 0
Teardown Count              : 0
-----
Last Bandwidth               : 290.4m
Last Requested Bandwidth     : 377.2m
Last Signaled Bandwidth      : 377.2m
Current Bandwidth            : 377.2m
Highest Bandwidth            : 377.2m
-----
Time for Next Sample request  : 59 seconds
Time for Next Adjustment     : 0 seconds
Time of Last Bandwidth Request : 2024 Jun 25 10:30:42
Time of Last Bandwidth Signal  : 2024 Jun 25 10:30:42
Time of Last Adjustment      : 2024 Jun 25 10:30:42
Time of Highest Bandwidth Marked : 2024 Jun 25 10:30:42
-----
Total Auto-Bandwidth Adjustments : 5
Successful Adjustments           : 5
Failed Adjustments               : 0
-----
Samples collected in the current adjustment cycle:
=====
Underflow:

Scenario 1 :
Current bandwidth is adjusted to 377.2 mbps. Then, send the sample rate which is less than underflow
threshold i.e, 317.3 mbps , 310 mbps.
As per the below output current bandwidth is more than underflow threshold bandwidth and
consecutively two samples are received and it is more than the underflow limit.
So current BW is adjusted to 317.3 mbps after 2 consecutive samples collected as per the Minimum
Average Bandwidth.

#show rsvp trunk auto-bandwidth detail
Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2203, Egress: 100.1.1.3
-----
Sample Interval               : 1 minutes
Adjustment Interval           : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth      : 200m
Minimum Bandwidth             : 100m
Maximum Bandwidth             : 500m
Overflow Threshold Bandwidth   : 40m
Underflow Threshold Bandwidth  : 30m
Overflow Threshold Activate Bandwidth : 0
Underflow Threshold Activate Bandwidth : 0
Overflow Limit                : 2
Underflow Limit               : 2
Max. Bandwidth Exceed Limit    : 2
-----
Max-BW-exceed-limit action    : teardown
Resignal-failure-action       : notify
Monitor Bandwidth             : No
-----
Minimum Average Bandwidth     : 317.3m
Maximum Average Bandwidth     : 0
Total Overflow Count          : 0

```

```

Consecutive Overflow Count      : 0
Consecutive Eligible Overflow Count : 0
Total Underflow Count          : 1
Consecutive Underflow Count     : 1
Consecutive Eligible Underflow Count : 1
Max. Bandwidth Exceed Count     : 0
Teardown Count                  : 0
-----

```

```

Last Bandwidth      : 290.4m
Last Requested Bandwidth : 377.2m
Last Signaled Bandwidth : 377.2m
Current Bandwidth    : 377.2m
Highest Bandwidth    : 377.2m
-----

```

```

Time for Next Sample request : 9 seconds
Time for Next Adjustment     : 3 minutes, 6 seconds
Time of Last Bandwidth Request : 2024 Jun 25 10:30:42
Time of Last Bandwidth Signal  : 2024 Jun 25 10:30:42
Time of Last Adjustment       : 2024 Jun 25 10:30:42
Time of Highest Bandwidth Marked : 2024 Jun 25 10:30:42
-----

```

```

Total Auto-Bandwidth Adjustments : 5
Successful Adjustments           : 5
Failed Adjustments                : 0
-----

```

```

Samples collected in the current adjustment cycle:
  [Sample 1-5]      : 317.3m

```

```
#show rsvp trunk auto-bandwidth detail
```

```
Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2204, Egress: 100.1.1.3
```

```

-----
Sample Interval      : 1 minutes
Adjustment Interval  : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth : 200m
Minimum Bandwidth    : 100m
Maximum Bandwidth    : 500m
Overflow Threshold Bandwidth : 40m
Underflow Threshold Bandwidth : 30m
Overflow Threshold Activate Bandwidth : 0
Underflow Threshold Activate Bandwidth : 0
Overflow Limit       : 2
Underflow Limit      : 2
Max. Bandwidth Exceed Limit : 2
-----

```

```

Max-BW-exceed-limit action : teardown
Resignal-failure-action    : notify
Monitor Bandwidth          : No
-----

```

```

Minimum Average Bandwidth : 0
Maximum Average Bandwidth : 0
Total Overflow Count      : 0
Consecutive Overflow Count : 0
Consecutive Eligible Overflow Count : 0
Total Underflow Count     : 0
Consecutive Underflow Count : 0
Consecutive Eligible Underflow Count : 0
Max. Bandwidth Exceed Count : 0
Teardown Count            : 0
-----

```

```

Last Bandwidth      : 377.2m
Last Requested Bandwidth : 317.3m
Last Signaled Bandwidth : 317.3m
Current Bandwidth    : 317.3m
Highest Bandwidth    : 377.2m
-----

```

```

Time for Next Sample request      : 56 seconds
Time for Next Adjustment         : 2 minutes, 46 seconds
Time of Last Bandwidth Request   : 2024 Jun 25 10:32:55
Time of Last Bandwidth Signal    : 2024 Jun 25 10:32:55
Time of Last Adjustment         : 2024 Jun 25 10:32:55
Time of Highest Bandwidth Marked : 2024 Jun 25 10:30:42

```

```

-----
Total Auto-Bandwidth Adjustments : 6
Successful Adjustments           : 6
Failed Adjustments               : 0
-----

```

Samples collected in the current adjustment cycle:

Scenario 2 :

Configure the Auto bandwidth Profile without configuring any underflow-limit. When all the samples in adjustment cycle receive with the underflow rate, then only underflow bandwidth adjustment will happen.

Below Example shows underflow limit as a Zero and current bandwidth is set to 8.5g and all the samples are received less then underflow-limit. So, the Bandwidth adjustment happens only after adjustment cycle.

#show rsvp trunk auto-bandwidth PE1-PE2

Session: PE1-PE2-Primary, Tunnel-id: 5002, LSP-ID: 2202, Egress: 100.1.1.3

```

-----
Sample Interval                  : 1 minutes
Adjustment Interval             : 5 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth       : 4g
Minimum Bandwidth               : 1g
Maximum Bandwidth               : 9g
Overflow Threshold Bandwidth    : 10% (851.2m)
Underflow Threshold Bandwidth  : 10% (851.2m)
Overflow Threshold Activate Bandwidth : 0
Underflow Threshold Activate Bandwidth : 0
Overflow Limit                  : 1
Underflow Limit                 : 0
Max. Bandwidth Exceed Limit    : 1

```

```

-----
Max-BW-exceed-limit action      : notify
Resignal-failure-action        : notify
Monitor Bandwidth               : No
-----

```

```

Minimum Average Bandwidth      : 6.5g
Maximum Average Bandwidth      : 0
Total Overflow Count           : 0
Consecutive Overflow Count     : 0
Consecutive Eligible Overflow Count : 0
Total Underflow Count         : 5
Consecutive Underflow Count    : 5
Consecutive Eligible Underflow Count : 5
Max. Bandwidth Exceed Count    : 0
Teardown Count                 : 0

```

```

-----
Last Bandwidth                  : 6.2g
Last Requested Bandwidth       : 8.5g
Last Signaled Bandwidth        : 8.5g
Current Bandwidth              : 8.5g
Highest Bandwidth              : 8.6g

```

```

-----
Time for Next Sample request    : 17 seconds
Time for Next Adjustment        : 0 seconds
Time of Last Bandwidth Request  : 2024 Jun 25 11:30:42
Time of Last Bandwidth Signal   : 2024 Jun 25 11:30:42
Time of Last Adjustment         : 2024 Jun 25 11:30:42
Time of Highest Bandwidth Marked : 2024 Jun 25 11:24:35

```

```

-----
Total Auto-Bandwidth Adjustments : 11
Successful Adjustments           : 11
Failed Adjustments               : 0
-----

```

Samples collected in the current adjustment cycle:
[Sample 1-5] : 6.5g 6.3g 6.4g 6.4g 6.4g

#show rsvp trunk auto-bandwidth

| -----+-----+-----+-----+-----+-----+-----+-----+-----+----- | | | | | | | | | |
|---|-------|------|------|-----------|----------|---------|---------|---------|--|
| -----+-----+-----+-----+-----+-----+-----+-----+-----+----- | | | | | | | | | |
| Trunk-Name | Trunk | LSP | Last | Requested | Signaled | Current | Highest | Adjust- | |
| Time Last-Adjust | ID | ID | BW | BW | BW | BW | BW | Left | |
| (sec) Time | | | | | | | | | |
| -----+-----+-----+-----+-----+-----+-----+-----+-----+----- | | | | | | | | | |
| PE1- | | | | | | | | | |
| PE2 | 5002 | 2202 | 8.5g | 6.5g | 6.5g | 6.5g | 8.6g | 277 | |
| 2024 Jun 25 11:30:42 | | | | | | | | | |

Boot on PE1

1. Create auto-bandwidth Profile.

(config)#router rsvp

2. Configure Auto bandwidth on boot and set the values for sample interval, Adjust interval and Adjust interval count.

(config-router)#auto-bandwidth-on-boot 1 5 1
(config-router)#commit

Validation

Verify auto bandwidth on boot adjustments information as below:

#show running-config rsvp
!
router rsvp
 auto-bandwidth-on-boot 1 5 1
!
!
!
!
!
!
!

#show rsvp trunk auto-bandwidth

*** On boot auto bandwidth is in progress for 2 minutes, 3 seconds ***

| -----+-----+-----+-----+-----+-----+-----+-----+-----+----- | | | | | | | | | |
|---|-------|------|------|-----------|----------|---------|---------|---------|----|
| -----+-----+-----+-----+-----+-----+-----+-----+-----+----- | | | | | | | | | |
| Trunk-Name | Trunk | LSP | Last | Requested | Signaled | Current | Highest | Adjust- | |
| Time Last-Adjust | ID | ID | BW | BW | BW | BW | BW | Left | |
| (sec) Time | | | | | | | | | |
| -----+-----+-----+-----+-----+-----+-----+-----+-----+----- | | | | | | | | | |
| PE-1_to_PE-2_ | | | | | | | | | |
| 1 | 5002 | 2202 | 0 | 0 | 0 | 200m | 144.6m | NA | NA |

Commands for RSVP Auto-Bandwidth

The RSVP auto-bandwidth uses the following configuration commands:

rsvp-auto-bandwidth

Use this command to configure an auto bandwidth profile. The profile will have default settings if any parameter not configured explicitly. User can configure parameters to their need within auto bandwidth profile.

Use `no` parameter of this command to delete auto bandwidth profile.

Command Syntax

```
rsvp-auto-bandwidth PROFILENAME
```

```
no rsvp-auto-bandwidth PROFILENAME
```

Parameters

<PROFILE_NAME>

Specifies the name assigned to the auto-bandwidth profile during configuration. The profile name can be a maximum of 64 characters in length.

Default

None

Command Mode

Config mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure an auto-bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#commit
(config-auto-bandwidth)#exit
(config)#
```

The following example describes how to delete the auto bandwidth profile:

```
#configure terminal
(config)#no rsvp-auto-bandwidth bwp
(config)#commit
```

sample-interval

Use this command to configure a sample interval value in minutes on the auto bandwidth profile. Sample interval determines the frequency at which rate samples collected from associated trunks. Sample interval must not be configured more than adjust interval as no samples can be collected within an adjustment cycle in such case.



Note: Sample interval timers run per auto bandwidth profile and not per associated trunks. So, in case of bandwidth adjustments on trunks before adjustment cycle completion will leave the newly formed session with less number of samples in the remaining part of adjustment cycle. In order to avoid very few samples being processed, minimum-samples command shall be configured in absolute or percentage format.

Use the `no` parameter to remove the sample interval configuration.

Command Syntax

```
sample-interval <1 - 10080>
```

```
no sample-interval
```

Parameters

<1-10080>

Specifies the sample interval value in minutes.

Default

5 minutes

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure sample interval in the auto bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#sample-interval 2
(config-auto-bandwidth)#commit
```

The following example describes how to remove configured sample interval in the auto bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no sample-interval
(config-auto-bandwidth)#commit
```

adjust-interval

Use this command to configure a adjust interval value in minutes on the auto bandwidth profile. Adjust interval determines the duration of the adjustment cycle. Bandwidth update decisions for active session of associated trunks are taken after every adjustment cycle. Adjust interval must not be configured less than sample interval as no samples can be collected within an adjustment cycle in such case.



Note: Adjust interval timers run per auto bandwidth profile and not per associated trunks. So, in case of bandwidth adjustments on trunks before adjustment cycle completion will leave the newly formed session with less number of samples in the remaining part of adjustment cycle. In order to avoid very few samples being processed, minimum-samples command shall be configured in absolute or percentage format.

Use the `no` parameter to remove the adjust interval configuration.

Command Syntax

```
adjust-interval <5 - 10080>
```

```
no adjust-interval
```

Parameters

<5-10080>

Specifies the adjust interval value in minutes.

Default

30 minutes

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure adjust interval in the auto bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#adjust-interval 60
(config-auto-bandwidth)#commit
```

The following example describes how to remove configured adjust interval in the auto bandwidth profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no adjust-interval
(config-auto-bandwidth)#commit
```


minimum-bandwidth

Use this command to configure minimum bandwidth on the auto bandwidth profile. Even when traffic flow is much lesser than minimum bandwidth, LSP will be reserved with the configured minimum bandwidth during bandwidth adjustment process. When auto bandwidth profile associated with trunk, LSP will be signaled with minimum bandwidth when initial bandwidth is not configured in the profile.

Use the `no` parameter to remove the minimum bandwidth configuration from the profile.

Command Syntax

```
minimum-bandwidth BANDWIDTH
```

```
no minimum-bandwidth
```

Parameters

BANDWIDTH

Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure minimum bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#minimum-bandwidth 100m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the minimum bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no minimum-bandwidth
(config-auto-bandwidth)#commit
```

maximum-bandwidth

Use this command to configure maximum bandwidth on the auto bandwidth profile. Even when traffic flow is much higher than maximum bandwidth, LSP will be reserved with the configured maximum bandwidth during bandwidth adjustment process. Operator notification is generated if the traffic rate samples collected are higher than the maximum bandwidth but the reservation is limited to maximum bandwidth.



Note: When maximum bandwidth is configured, even a single traffic rate sample crossing the maximum bandwidth will trigger an MBB with maximum bandwidth reserved. If user doesn't wish to trigger an MBB for single sample of maximum bandwidth exceed, maximum-bandwidth-exceed-limit shall be configured with a value to mention the number of consecutive samples to cross maximum bandwidth to take further action.

Use the `no` parameter to remove the maximum bandwidth configuration from the profile.

Command Syntax

```
maximum-bandwidth BANDWIDTH
```

```
no maximum-bandwidth
```

Parameters

BANDWIDTH

Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure maximum bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#maximum-bandwidth 900m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the maximum bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no maximum-bandwidth
(config-auto-bandwidth)#commit
```

initial-bandwidth

Use this command to configure initial bandwidth on the auto bandwidth profile. When auto bandwidth profile associated with trunk, LSP will be signalled with initial bandwidth when initial bandwidth is configured in the profile. For trunks which are already associated with auto bandwidth profile and the system going through reload, initial bandwidth will not be applicable as on boot computation will trigger to update active sessions with bandwidth as per the on boot period traffic rate sample computation.

Use the `no` parameter to remove the initial bandwidth configuration from the profile.

Command Syntax

```
initial-bandwidth BANDWIDTH
```

```
no initial-bandwidth
```

Parameters

BANDWIDTH

Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure initial bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#initial-bandwidth 500m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the initial bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no initial-bandwidth
(config-auto-bandwidth)#commit
```

underflow-threshold

Use this command to configure underflow threshold in percentage or absolute value format on the auto bandwidth profile. Underflow threshold sets the amount of reduction in traffic rate sample to detect an eligible underflow. As an example, absolute underflow threshold 10m when current bandwidth is 200m means, a traffic rate sample of 190.1m will not be considered eligible underflow sample and a sample of 189.9m will be considered eligible underflow sample.

When all the traffic rate samples collected for an auto bandwidth profile associated trunk cross underflow threshold in an adjustment cycle, then the highest eligible traffic rate sample will be considered to re-signal the session with new bandwidth at the end of an adjustment cycle.

Constraints like underflow-limit and underflow-threshold-activate-bandwidth will add additional logic on how bandwidth update action is taken. This will be discussed in respective sections.

If underflow threshold is not configured, then minor reduction in traffic rate sample also will be considered as eligible underflow bandwidth sample. So, underflow and overflow threshold is a recommended configuration even though it is not mandatory.

When underflow threshold is configured in percentage, the threshold will be computed based on the current bandwidth and the percentage value. Example, underflow threshold 10% for a current bandwidth of 100m means a sample of 90m or lesser will be considered eligible underflow sample. Underflow threshold can be configured either as absolute value or in percentage but not both.

Use the `no` parameter to remove the underflow bandwidth configuration from the profile.

Command Syntax

```
underflow-threshold (percent <1-100>) | (absolute BANDWIDTH)
```

```
no underflow-threshold (percent | absolute)
```

Parameters

<1-100>

Specifies the underflow threshold value in percentage.

BANDWIDTH

Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure underflow bandwidth in percentage format in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
```

```
(config-auto-bandwidth)#underflow-threshold percent 10
(config-auto-bandwidth)#commit
```

The following example describes how to remove the underflow bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no underflow-threshold percent
(config-auto-bandwidth)#commit
```

overflow-threshold

Use this command to configure overflow threshold in percentage or absolute value format on the auto bandwidth profile. Overflow threshold sets the amount of increase in traffic rate sample required to detect an eligible overflow. As an example, absolute overflow threshold 10m when current bandwidth is 200m means, a traffic rate sample of 209.9m will not be considered eligible overflow sample and a sample of 210.1m will be considered eligible overflow sample.

When a traffic rate sample collected for an auto bandwidth profile associated trunk crosses overflow threshold in an adjustment cycle, then the highest eligible traffic rate sample will be considered to re-signal the session with new bandwidth at the end of adjustment cycle.

Constraints like overflow-limit and overflow-threshold-activate-bandwidth will add additional logic on how bandwidth update action is taken. This will be discussed in respective sections.

If overflow threshold is not configured, then minor increase in traffic rate sample also will be considered as eligible overflow bandwidth sample. So, underflow and overflow threshold is a recommended configuration even though it is not mandatory.

When overflow threshold is configured in percentage, the threshold will be computed based on the current bandwidth and the percentage value. Example, overflow threshold 10% for a current bandwidth of 100m means a sample of 110m or more will be considered eligible overflow sample. Overflow threshold can be configured either as absolute value or in percentage but not both.

Use the `no` parameter to remove the underflow bandwidth configuration from the profile.

Command Syntax

```
overflow-threshold (percent <1-100>) | (absolute BANDWIDTH)
```

```
no overflow-threshold (percent | absolute)
```

Parameters

<1-100>

Specifies the overflow threshold value in percentage.

BANDWIDTH

Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure underflow bandwidth in absolute format in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
```

```
(config-auto-bandwidth)#overflow-threshold absolute 10m  
(config-auto-bandwidth)#commit
```

The following example describes how to remove the underflow bandwidth configuration from the profile:

```
#configure terminal  
(config)#rsvp-auto-bandwidth bwp  
(config-auto-bandwidth)#no overflow-threshold absolute  
(config-auto-bandwidth)#commit
```

underflow-threshold-activate-bandwidth

Use this command to configure absolute bandwidth range to allow bandwidth re-signalling when underflow threshold and underflow limit criteria matched. This configuration helps to limit the underflow bandwidth reservation update for certain range of bandwidth.

As an example, if the current bandwidth is 500m and the underflow threshold is 10%. So, normally, if all traffic rate samples collected are in the range of 400m to 450m, session will be re-signalled to reserve new bandwidth. However, if underflow-threshold-activate-bandwidth is configured as 300m, then the traffic rate samples in the range of 400m to 450m will not trigger bandwidth update. Only when the traffic rate samples are less than 300m, then it will be considered as eligible sample.

The configuration creates an absolute bandwidth range for underflow samples to be eligible. The bandwidth range for underflow eligibility will be minimum bandwidth (or zero when minimum bandwidth is not configured) to underflow-threshold-activate-bandwidth value. When this command is not configured, there won't be any such absolute range and only underflow-threshold and underflow-limit will be considered for computation, if configured.

Use the `no` parameter to remove the underflow threshold activate bandwidth configuration from the profile.

Command Syntax

```
underflow-threshold-activate-bandwidth BANDWIDTH
```

```
no underflow-threshold-activate-bandwidth
```

Parameters

BANDWIDTH

Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure underflow threshold activate bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#underflow-threshold-activate-bandwidth 500m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the underflow threshold activate bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no underflow-threshold-activate-bandwidth
(config-auto-bandwidth)#commit
```


overflow-threshold-activate-bandwidth

Use this command to configure absolute bandwidth range to allow bandwidth re-signalling when overflow threshold and overflow limit criteria matched. This configuration helps to limit the overflow bandwidth reservation update for certain range of bandwidth.

As an example, if the current bandwidth is 100m and the overflow threshold is 10%. Normally, if a traffic rate sample collected is in the range of 110m to 150m, session will be re-signalled to reserve new bandwidth. However, if overflow-threshold-activate-bandwidth is configured as 300m, then the traffic rate samples in the range of 110m to 150m will not trigger bandwidth update. Only when the traffic rate samples are more than 300m, then it will be considered as eligible sample.

The configuration creates an absolute bandwidth range for overflow samples to be eligible. The bandwidth range for overflow eligibility will be overflow-threshold-activate-bandwidth value to a practical infinity. When this command is not configured, there won't be any such absolute range and only overflow-threshold and overflow-limit will be considered for computation, if configured.

Use the `no` parameter to remove the overflow threshold activate bandwidth configuration from the profile.

Command Syntax

```
overflow-threshold-activate-bandwidth BANDWIDTH
```

```
no overflow-threshold-activate-bandwidth
```

Parameters

BANDWIDTH

Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure overflow threshold activate bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#overflow-threshold-activate-bandwidth 500m
(config-auto-bandwidth)#commit
```

The following example describes how to remove the overflow threshold activate bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no overflow-threshold-activate-bandwidth
(config-auto-bandwidth)#commit
```

underflow-limit

Use this command to configure underflow limit on the auto bandwidth profile. When underflow limit is configured, if the traffic rate samples collected on the associated session consecutively crosses underflow threshold for underflow limit times, then the bandwidth adjustment will be triggered immediately without waiting for adjustment cycle completion. When underflow-threshold-activate-bandwidth is configured, even this criteria is considered to mark a sample as eligible underflow sample.

Only when underflow limit is configured, underflow adjustment may happen before the completion of adjustment cycle. Otherwise, underflow adjustment considered only at the completion of adjustment cycle when all samples found to be eligible underflow sample.

Use the `no` parameter to remove the underflow limit configuration from the profile.

Command Syntax

```
underflow-limit <1-10080>
```

```
no underflow-limit
```

Parameters

<1-10080>

Specifies the underflow limit value for consecutive eligible samples.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure underflow limit in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#underflow-limit 3
(config-auto-bandwidth)#commit
```

The following example describes how to remove the underflow limit configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no underflow-limit
(config-auto-bandwidth)#commit
```

overflow-limit

Use this command to configure overflow limit on the auto bandwidth profile. When overflow limit is configured, if the traffic rate samples collected on the associated session consecutively crosses overflow threshold for overflow limit times, then the bandwidth adjustment will be triggered immediately without waiting for adjustment cycle completion. When overflow-threshold-activate-bandwidth is configured, even this criteria is considered to mark a sample as eligible underflow sample.

Only when overflow limit is configured, overflow adjustment may happen before the completion of adjustment cycle. Otherwise, overflow adjustment considered only at the completion of adjustment cycle when a sample found to be eligible overflow sample.

If the traffic rate sample crosses maximum bandwidth, then maximum-bandwidth-exceed-limit configuration comes into picture and by default, a single sample crossing maximum bandwidth triggers bandwidth update. This situation is different from overflow scenario.

Use the `no` parameter to remove the overflow limit configuration from the profile.

Command Syntax

```
overflow-limit <1-10080>
```

```
no overflow-limit
```

Parameters

<1-10080>

Specifies the overflow limit value for consecutive eligible samples.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure overflow limit in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#overflow-limit 3
(config-auto-bandwidth)#commit
```

The following example describes how to remove the overflow limit configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no overflow-limit
(config-auto-bandwidth)#commit
```

maximum-bandwidth-exceed-limit

Use this command to configure maximum bandwidth exceed limit on the auto bandwidth profile. When maximum bandwidth exceed limit is configured, if the traffic rate samples collected on the associated session consecutively crosses maximum bandwidth for maximum-bandwidth-exceed-limit times, then the action will be triggered immediately without waiting for adjustment cycle completion. When maximum-bandwidth-exceed-limit is not configured, a single sample exceeding maximum bandwidth will trigger an action which is re-signal with updated bandwidth or restart the session with initial or minimum bandwidth based on the action configured.

When maximum bandwidth is not configured, maximum bandwidth exceed limit configuration doesn't have any significance. Overflow limit and maximum bandwidth exceed limits are independent commands with different significance with latter associated with maximum bandwidth.

Use the `no` parameter to remove the maximum bandwidth exceed limit configuration from the profile.

Command Syntax

```
maximum-bandwidth-exceed-limit <1-10080>
```

```
no maximum-bandwidth-exceed-limit
```

Parameters

<1-10080>

Specifies the maximum bandwidth exceed limit value for consecutive eligible samples.

Default

1

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure maximum bandwidth exceed limit in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#maximum-bandwidth-exceed-limit 2
(config-auto-bandwidth)#commit
```

The following example describes how to remove the maximum bandwidth exceed limit configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no maximum-bandwidth-exceed-limit
(config-auto-bandwidth)#commit
```

maximum-bandwidth-exceed-action

Use this command to configure maximum bandwidth exceed action on the auto bandwidth profile. When the traffic rate samples collected on the associated session consecutively crosses maximum bandwidth for maximum-bandwidth-exceed-limit times (or one time if limit is not configured), then the action to be triggered will be decided by this configuration. If not configured, default action is to re-signal the session with maximum bandwidth or ignore if session is already signalled with maximum bandwidth. In any case, user will be notified about the maximum bandwidth being exceeded. However, with exceed action configured as teardown, session will be released and restarted with initial bandwidth or minimum bandwidth if initial bandwidth is not configured.

This action will lead to service interruption if there are no alternate transport. So, this configuration is recommended to be used with full awareness of the impact.

Use the `no` parameter to remove the maximum bandwidth exceed action configuration from the profile.

Command Syntax

```
maximum-bandwidth-exceed-action (teardown)
```

```
no maximum-bandwidth-exceed-action
```

Parameters

teardown

Teardown the session exceeding maximum bandwidth.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure maximum bandwidth exceed action in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#maximum-bandwidth-exceed-action teardown
(config-auto-bandwidth)#commit
```

The following example describes how to remove the maximum bandwidth exceed action configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no maximum-bandwidth-exceed-action
(config-auto-bandwidth)#commit
```

resignal-failure-action

Use this command to configure an action on the auto bandwidth profile when the bandwidth update re-signalling fails on the associated session. By default, if re-signalling fails (3 attempts) for the updated bandwidth, it will be noted down as re-signalling failure and session will continue with its current bandwidth reservation. If severe actions to be taken on such re-signal failure, then teardown action can be configured which will release the current session and restart freshly with initial bandwidth or minimum bandwidth when initial bandwidth is not configured.

This action will lead to service interruption if there are no alternate transport. So, this configuration is recommended to be used with full awareness of the impact.

Use the `no` parameter to remove the re-signal failure action configuration from the profile.

Command Syntax

```
resignal-failure-action (teardown)
```

```
no signal-failure-action
```

Parameters

teardown

Specifies the teardown the session when re-signalling with new bandwidth fails.

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure re-signal failure action in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#resignal-failure-action teardown
(config-auto-bandwidth)#commit
```

The following example describes how to remove the re-signal failure action configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no signal-failure-action
(config-auto-bandwidth)#commit
```

sync-bandwidth

Use this command to configure bandwidth synchronization for primary and secondary sessions of an auto bandwidth profile associated trunk. With this configuration, in case the associated trunk is configured with primary and secondary sessions, every time primary session goes through a bandwidth update, secondary session also will be re-signalled with primary session's bandwidth. Thus, secondary path is determined with proper reservation constraints to ensure it is in the correct bandwidth reserved state when traffic switches to secondary.

Use the `no` parameter to remove synchronise bandwidth configuration from the profile.

Command Syntax

```
sync-bandwidth
```

```
no sync-bandwidth
```

Parameters

None

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure synchronize bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#sync-bandwidth
(config-auto-bandwidth)#commit
```

The following example describes how to remove the synchronise bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no sync-bandwidth
(config-auto-bandwidth)#commit
```

monitor-bandwidth

Use this command to configure only monitor the traffic rate samples and computation without taking any action. This command can be used to monitor the traffic behaviour without updating the active sessions. With this configuration, in case of overflow, underflow, adjustment cycle completion time computation results, maximum bandwidth exceed, etc., notification is provided without taking any action.

Use the `no` parameter to remove monitor bandwidth configuration from the profile.

Command Syntax

```
monitor-bandwidth
```

```
no monitor-bandwidth
```

Parameters

None

Default

None

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure monitor bandwidth in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#monitor-bandwidth
(config-auto-bandwidth)#commit
```

The following example describes how to remove the monitor bandwidth configuration from the profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no monitor-bandwidth
(config-auto-bandwidth)#commit
```


minimum-samples

Use this command to configure the minimum samples required in an adjustment cycle for bandwidth processing. Sample timers and Adjust timers are executed per auto bandwidth profile and not per associated trunk. Thus, there are scenarios of a trunk going through a bandwidth update few minutes ago and again ends up with adjustment cycle completion processing with very few samples collected. In order to avoid such scenarios, minimum samples required in an adjustment cycle to process the bandwidth shall be configured.

Configuration is accepted in both absolute and in percentage format. This gives user the flexibility to choose the format that suites their need. If sample interval and adjust interval expected to be fixed, then absolute configuration helps providing the requirement of exact number of minimum samples required to process. If exact number isn't important and there are chances of changing adjust interval or sample interval in future, then percentage format can be chosen. However, only one of the formats can be configured.

By default, even if there is one traffic rate sample during adjustment cycle completion, bandwidth will be processed. So, it will be recommended to have this configuration if users are keen on minimum of certain samples to be considered for bandwidth computation.

Use the `no` parameter to remove the minimum sample configuration from the profile.

Command Syntax

```
minimum-samples (percent <1-100>) | (absolute <1-10080>)
```

```
no underflow-limit (percent | absolute)
```

Parameters

<1-10080>

Specifies the absolute value for minimum samples required in an adjustment cycle.

<1-100>

Specifies the minimum sample percentage required in an adjustment cycle.

Default

1 sample

Command Mode

Auto bandwidth mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example describes how to configure minimum samples in percentage format in a profile:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#minimum-samples percent 70
(config-auto-bandwidth)#commit
```

The following example describes how to remove the minimum samples configuration from the profile:

```
#configure terminal
```

```
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#no minimum-samples percent
(config-auto-bandwidth)#commit
```

auto-bandwidth

Use this command to attach an auto bandwidth profile to a trunk. When the auto bandwidth profile is attached to the trunk, active session will be re-signalled with initial bandwidth configured in the auto bandwidth profile or minimum bandwidth configured if initial bandwidth is not configured. Bandwidth update will be triggered only if there is variation in the bandwidth to be initialized. Attaching or detaching an auto bandwidth profile doesn't trigger any session flap and doesn't cause traffic impact.

When an auto bandwidth profile is associated with first trunk, sample interval and adjust interval timers will start and are stopped when the profile is removed from the last trunk.

Manual bandwidth configuration for the sessions and auto bandwidth profile attach are mutually exclusive and the configuring both of them on a trunk is not allowed.

Use the `no` parameter to remove the auto bandwidth profile from the trunk.

Command Syntax

```
auto-bandwidth PROFILENAME
```

```
no auto-bandwidth PROFILENAME
```

Parameters

PROFILENAME

Specifies the name of the auto bandwidth profile.

Default

None

Command Mode

Trunk mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to associate an auto bandwidth profile to a trunk:

```
#configure terminal
(config)#rsvp-auto-bandwidth bwp
(config-auto-bandwidth)#exit
(config)#rsvp-trunk t1
(config-trunk)#auto-bandwidth bwp
(config-trunk)#commit
```

The following example describes how to remove the auto bandwidth profile from the trunk:

```
#configure terminal
(config)#rsvp-trunk t1
(config-trunk)#no auto-bandwidth bwp
(config-trunk)#commit
```

auto-bandwidth-on-boot

Use this command to configure on boot sample interval, adjust interval and number of adjustment cycles. When the system is reloaded and comes up, all active sessions of trunks associated with auto bandwidth profiles run a relatively faster adjustment cycle with quicker sample collection to settle the sessions with accurate bandwidth reservation.

By default, sample interval is 1 minute, adjust interval is 5 minutes and the adjustment cycle runs one time. After the adjustment cycle completion, samples of each associated trunks computed to re-signal the sessions with updated bandwidth. Then the auto bandwidth profile based adjustment cycle starts. If user wishes to run the boot up time rigorous sample computation for longer duration or multiple rounds, then it shall be configured. The configurations will apply from system reload if the configuration is saved. Properties of auto bandwidth profiles will not be applied during boot up time computation.

Use the `no` parameter to remove the auto bandwidth profile from the trunk.

Command Syntax

```
auto-bandwidth-on-boot <1-10080> <1-10080> <1-10>
```

```
no auto-bandwidth-on-boot
```

Parameters

<1-10080>

On boot sample interval value in minutes.

<1-10080>

On boot adjustment interval value in minutes.

<1-10>

Specifies the number of adjustment cycles to run on boot.

Default

Sample interval 1 minute, adjust interval 5 minutes and 1 adjustment cycle.

Command Mode

Router mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to configure on boot auto bandwidth parameters:

```
#configure terminal
(config)#router rsvp
(config-router)#auto-bandwidth-on-boot 1 10 3
(config-router)#commit
```

The following example describes how to reset on boot auto bandwidth parameters:

```
#configure terminal
(config)#router rsvp
(config-router)#no auto-bandwidth-on-boot
(config-router)#commit
```

force-auto-bandwidth-adjustment

Use this command to force a bandwidth adjustment on a trunk associated with auto bandwidth profile. When the command is executed without bandwidth value mentioned, traffic rate samples collected till the time are used to compute the bandwidth to be adjusted. In case of bandwidth value mentioned in the command, the bandwidth is verified for eligibility and bandwidth update will be triggered.

Command Syntax

```
rsvp-trunk TRUNKNAME force-auto-bandwidth-adjustment (BANDWIDTH|)
```

Parameters

TRUNKNAME

Specifies the name of the trunk to go through forced bandwidth adjustment.

BANDWIDTH

Specifies the bandwidth value in the range of 1k to 999g.

Default

None

Command Mode

Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how you can force a bandwidth adjustment for a trunk with an auto bandwidth profile:

```
#rsvp-trunk t1 force-auto-bandwidth-adjustment
```

clear rsvp auto-bandwidth

Use this command to reset the auto bandwidth adjustment cycle by clearing all the traffic samples collected by the associated trunks and by restarting sample and adjust timers. If auto bandwidth profile name is not mentioned, then all trunks associated with any auto bandwidth profile will be reset and computation will start freshly.

Command Syntax

```
clear rsvp auto-bandwidth (PROFILENAME|)
```

Parameters

PROFILENAME

Specifies the name of the auto bandwidth profile.

Default

None

Command Mode

Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to restart processing of an auto bandwidth profile:

```
#clear rsvp auto-bandwidth bwp
```

s

clear rsvp trunk auto-bandwidth-statistics

Use this command to clear the statistics maintained on a trunk associated with auto bandwidth profile. Statistics will be mainly the highest watermarked bandwidth, last adjusted bandwidth, how many times adjustment triggered, status of the adjustment trigger, etc. This command will only clear the auto bandwidth statistics for the trunk and doesn't impact the operation of auto bandwidth including the traffic rate samples collected for the current adjustment cycle.

Command Syntax

```
clear rsvp trunk TRUNKNAME auto-bandwidth-statistics
```

Parameters

TRUNKNAME

Specifies the name of the trunk associated with auto bandwidth profile.

Default

None

Command Mode

Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example describes how to clear the auto bandwidth statistics on a trunk which is associated with the auto bandwidth profile:

```
#clear rsvp trunk t1 auto-bandwidth-statistics
```

show rsvp auto-bandwidth

Use this command to display auto bandwidth profile specific information.

Command Syntax

```
show rsvp auto-bandwidth
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for viewing an auto bandwidth summary of all the trunks associated with auto bandwidth profile:

```
#show rsvp auto-bandwidth
Profile Name : bwp
-----
Sample Interval : 5 minutes (due in 4 minutes)
Adjust Interval : 30 minutes (due in 29 minutes)
Minimum Samples required for processing : 1
Initial Bandwidth : 0
Minimum bandwidth : 0
Maximum bandwidth : 100m
Underflow Threshold Bandwidth : 5m
Overflow Threshold Bandwidth : 5m
Underflow Threshold Activate Bandwidth : 0
Overflow Threshold Activate Bandwidth : 0
Overflow Limit : 3
Underflow Limit : 3
Maximum Bandwidth Exceed Limit : 1
Maximum Bandwidth Exceed Action : notify
Re-signal Failure Action : notify
Sync Bandwidth : No
Monitor Bandwidth : No
No. of trunks associated : 1
```


show rsvp auto-bandwidth detail

Use this command to display a specific auto bandwidth profile information or all auto bandwidth profile information along with associated trunk details.

Command Syntax

```
show rsvp auto-bandwidth detail
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example is for viewing all the auto bandwidth profiles along with associated trunks:

```
#show rsvp auto-bandwidth detail
```

```
Profile Name : bwp
```

```
-----
Sample Interval           : 5 minutes (due in 4 minutes)
Adjust Interval           : 30 minutes (due in 29 minutes)
Minimum Samples required for processing : 1
Initial Bandwidth         : 0
Minimum bandwidth         : 0
Maximum bandwidth         : 100m
Underflow Threshold Bandwidth : 5m
Overflow Threshold Bandwidth : 5m
Underflow Threshold Activate Bandwidth : 0
Overflow Threshold Activate Bandwidth : 0
Overflow Limit            : 3
Underflow Limit           : 3
Maximum Bandwidth Exceed Limit : 1
Maximum Bandwidth Exceed Action : notify
Re-signal Failure Action   : notify
Sync Bandwidth            : No
Monitor Bandwidth         : No
No. of trunks associated   : 1
```

```
-----+-----+-----+-----+-----+-----+-----+-----+-----
Trunk-Name  Trunk   LSP    Last   Requested  Signaled  Current  Highest  LastAdjust
            ID      ID      BW      BW         BW         BW         BW         Time
-----+-----+-----+-----+-----+-----+-----+-----+-----
t1          5001    2201    10.1m  22.5m      22.5m     22.5m     35.5m     2024 Jul 23
```

show rsvp auto-bypass-group

Use this command to display the mapping of the group ID value used in the auto bypass name to exclude interface index and SRLG values. This command helps understand the value present in the auto bypass tunnel names.

Command Syntax

```
show rsvp auto-bypass-group
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

Example for viewing all the auto bandwidth profiles:

```
#show rsvp auto-bypass-group
SRLG-disjoint Configured: Forced
Group-ID  Refcnt    Exclude-IfIndex  SRLG Values
-----+-----+-----+-----
100       3           10001            10 20
```

show rsvp trunk auto-bandwidth

Use this command to display the information of all the trunks associated with the auto bandwidth profile. This show command will display high level information like what is the last bandwidth, current bandwidth, last adjustment time, time left in adjustment cycle in seconds, etc.

Command Syntax

```
show rsvp trunk auto-bandwidth
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for viewing an auto bandwidth summary osf all the trunks associated with auto bandwidth profile:

```
#show rsvp trunk auto-bandwidth
```

| Trunk-Name | Trunk ID | LSP ID | Last BW | Requested BW | Signaled BW | Current BW | Highest BW | Adjust-Time Left (sec) | Last-Adjust Time |
|------------|----------|--------|---------|--------------|-------------|------------|------------|------------------------|------------------|
| t1 | 5001 | 2201 | 10.1m | 22.5m | 22.5m | 22.5m | 35.5m | 1142 | 2024 Jul 23 |

show rsvp trunk auto-bandwidth detail

Use this command to display the information of a trunk or all the trunks associated with the auto bandwidth profile. This command will provide detailed information of the auto bandwidth related statistics on the trunk as well as details of traffic rate samples collected in an adjust cycle and the time left for next sample collection, etc.

Command Syntax

```
show rsvp trunk auto-bandwidth detail
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

Example for viewing the auto bandwidth details of all the trunks associated with auto bandwidth profile:

```
#show rsvp trunk auto-bandwidth detail
:Session: tl-Primary, Tunnel-id: 5001, LSP-ID: 2201, Egress: 2.2.2.2
```

```
-----
Sample Interval           : 5 minutes
Adjustment Interval      : 30 minutes
Minimum Samples required for processing : 1
Initialization Bandwidth : 0
Minimum Bandwidth        : 0
Maximum Bandwidth        : 100m
Overflow Threshold Bandwidth : 5m
Underflow Threshold Bandwidth : 5m
Overflow Threshold Activate Bandwidth : 0
Underflow Threshold Activate Bandwidth : 0
Overflow Limit           : 3
Underflow Limit          : 3
Max. Bandwidth Exceed Limit : N/A
-----
```

```
Max-BW-exceed-limit action : notify
Resignal-failure-action    : notify
Monitor Bandwidth          : No
-----
```

```
Minimum Average Bandwidth : 0
Maximum Average Bandwidth : 22.5m
Total Overflow Count      : 1
Consecutive Overflow Count : 1
Consecutive Eligible Overflow Count : 1
Total Underflow Count     : 0
Consecutive Underflow Count : 0
Consecutive Eligible Underflow Count : 0
Max. Bandwidth Exceed Count : 0
Teardown Count           : 0
-----
```

```
Last Bandwidth           : 10.2m
Last Requested Bandwidth : 15.6m
Last Signaled Bandwidth  : 15.6m
Current Bandwidth        : 15.6m
```

```
Highest Bandwidth                : 35.3m
-----
Time for Next Sample request      : 1 minutes, 20 seconds
Time for Next Adjustment         : 16 minutes, 30 seconds
Time of Last Bandwidth Request   : 2024 Jul 23 11:32:44
Time of Last Bandwidth Signal    : 2024 Jul 23 11:32:44
Time of Last Adjustment         : 2024 Jul 23 11:32:44
Time of Highest Bandwidth Marked : 2024 Jul 23 11:14:37
-----
Total Auto-Bandwidth Adjustments : 2
Successful Adjustments           : 2
Failed Adjustments               : 0
-----
Samples collected in the current adjustment cycle:
[Samples 1-5]      : 17.5m      18.3m      22.5m
```

RSVP LSP Reoptimization

Overview

RSVP session re-optimization allows the system to periodically recalculate Label Switched Paths (LSPs) based on topology changes to ensure optimal path utilization. Re-optimization isn't enabled on the LSPs by default except detour LSPs and are configurable by the user. Re-optimization enabled LSPs are added to a reoptimization wheel timer in CSPF server with an offset computed based on the LSP_ID. LSP details in the CSPF server are Enhanced to provide operational visibility into the re-optimization lifecycle by displaying the "remaining time" until the next re-optimization trigger for each LSP. This visibility is essential for troubleshooting instances where topology variations occur but re-optimization is not immediately observed, allowing users to verify the status of the re-optimization wheel timer. In case the next reoptimization going to take time but an immediate reoptimization intended, then the force reoptimize utility allows to trigger the reoptimization per session.

Characteristics

- Re-optimization applies to primary, secondary, auto-bypass, and manual-bypass LSPs on a configuration basis.
- Re-optimization for detour sessions are supported by default.
- Adds the "remaining time" field to indicate when the next re-optimization attempt will occur for a given LSP.
- Force re-optimization support extended for all type of LSPs.

Benefits

- The re-optimization configuration allows users to decide if an LSP should be recalculated, effectively reducing processing overhead by preventing unnecessary updates.
- Displaying reop-remaining-time improves troubleshooting and monitoring by exposing real-time timer information.
- Displaying LSP name and state allows users to correlate CSPF re-optimization activity with RSVP session states.

Limitations

- Re-optimization will be active on the LSPs which are in established state.
- The remaining time cannot be displayed under show rsvp session details since the timer is managed by the CSPF server (OSPF/IS-IS), not by RSVP.

Prerequisites

- RSVP re-optimization must be enabled for the LSP in one of the following modes:
 - trunk (for primary/secondary LSPs)
 - bypass-trunk (for manual bypass)
 - auto-bypass (for automatic bypass)
- CSPF-based path computation must be operational.


- Applicable to RSVP-TE deployments using CSPF for LSP computation.
- To capture RSVP session re-optimization events, set the logging level to notify (level 4).

Configuration

Topology

The provided topology illustrates an MPLS network core interconnecting two customer edge (CE) devices through a series of provider edge (PE) and provider (P) routers.

Figure 36. RSVP-TE Core Network Topology for LSP Re-optimization Monitoring



P1

1. Enable RSVP.

```
#configure terminal
(config)#router rsvp
(config-router)#commit
(config-router)#exit
```

2. Configure the interfaces.

```
(config)#interface lo
(config-if)#ip address 30.0.1.2/32 secondary
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe4
(config-if)#load-interval 30
(config-if)#ip address 10.0.1.4/31
(config-if)#mtu 9216
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe44
(config-if)#load-interval 30
```

```
(config-if)#ip address 10.0.1.3/31
(config-if)#mtu 9216
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

3. Configure OSPF process.

```
(config)#router ospf 100
(config-router)#ospf router-id 30.0.1.2
(config-router)#network 10.0.1.2/31 area 0.0.0.0
(config-router)#network 10.0.1.4/31 area 0.0.0.0
(config-router)#network 30.0.1.2/32 area 0.0.0.0
(config-router)#commit
(config-router)#exit
```

P2

1. Enable RSVP.

```
#configure terminal
(config)#router rsvp
(config-router)#commit
(config-router)#exit
```

2. Configure the interfaces.

```
(config)#interface ce49
(config-if)#load-interval 30
(config-if)#ip address 10.0.1.27/31
(config-if)#mtu 9216
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

```
(config)#interface lo
(config-if)#ip address 30.0.1.5/32 secondary
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe26
(config-if)#load-interval 30
(config-if)#mtu 9600
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe26.20
(config-if)#encapsulation dot1q 20
(config-if)#load-interval 30
(config-if)#ip address 10.0.1.21/31
(config-if)#mtu 9216
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

3. Configure OSPF process.

```
(config)#router ospf 100
(config-router)#ospf router-id 30.0.1.5
(config-router)#network 10.0.1.20/31 area 0.0.0.0
(config-router)#network 10.0.1.26/31 area 0.0.0.0
```



```
(config-router)#network 30.0.1.5/32 area 0.0.0.0
(config-router)#commit
(config-router)#exit
```

PE1

1. Configure VRF.

```
#configure terminal
(config)#ip vrf VRF51
(config-vrf)#rd 1.1.1.1:51
(config-vrf)#route-target export 100:51
(config-vrf)#route-target import 300:51
(config-vrf)#commit
(config-vrf)#exit
```

2. Enable RSVP.

```
(config)#router rsvp
(config-router)#commit
(config-router)#exit
```

3. Configure the interfaces.

```
(config)#interface lo
(config-if)#ip address 30.0.1.1/32 secondary
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe24
(config-if)#load-interval 30
(config-if)#ip address 10.0.1.2/31
(config-if)#mtu 9216
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe25
(config-if)#load-interval 30
(config-if)#mtu 9600
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe25.6
(config-if)#encapsulation dot1q 6
(config-if)#load-interval 30
(config-if)#ip address 10.0.1.6/31
(config-if)#mtu 9216
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe26
(config-if)#load-interval 30
(config-if)#mtu 9600
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe26.20
(config-if)#encapsulation dot1q 20
(config-if)#load-interval 30
(config-if)#ip address 10.0.1.20/31
(config-if)#mtu 9216
```

```
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe27
(config)#load-interval 30
(config)#mtu 9600
(config)#commit
(config)#exit
```

```
(config)#interface xe27.51
(config)#description L3VPN-VRF51
(config)#encapsulation dot1q 51
(config)#load-interval 30
(config)#ip vrf forwarding VRF51
(config)#ip address 161.0.51.1/24
(config)#mtu 9216
(config)#commit
(config)#exit
```

4. Configure OSPF process.

```
(config)#router ospf 100
(config-router)#ospf router-id 30.0.1.1
(config-router)#network 10.0.1.2/31 area 0.0.0.0
(config-router)#network 10.0.1.6/31 area 0.0.0.0
(config-router)#network 10.0.1.20/31 area 0.0.0.0
(config-router)#network 30.0.1.1/32 area 0.0.0.0
(config-router)#commit
(config-router)#exit
```

5. Configure the BGP.

```
(config)#router bgp 65535
(config-router)#neighbor 30.0.1.3 remote-as 65535
(config-router)#neighbor 30.0.1.3 update-source lo
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 30.0.1.3 activate
(config-router-af)#exit-address-family
(config-router-af)# address-family vpnv4 unicast
(config-router-af)#neighbor 30.0.1.3 activate
(config-router-af)#exit-address-family
(config-router-af)#address-family ipv4 vrf VRF51
(config-router-af)#redistribute connected
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

6. Configure RSVP trunk and enable reoptimization.

```
(config)#rsvp-trunk PE1-PE2 ipv4
(config-trunk)#reoptimize
(config-trunk)#to 30.0.1.3
(config-trunk)#commit
(config-trunk)#exit
```

PE2

1. Configure VRF.

```
#configure terminal
(config)#ip vrf VRF51
(config-vrf)#rd 3.3.3.3:51
(config-vrf)#route-target import 100:51
(config-vrf)#route-target export 300:51
```

```
(config-vrf)#commit
(config-vrf)#exit
```

2. Enable RSVP.

```
(config)#router rsvp
(config)#commit
(config)#exit
```

3. Configure the interfaces.

```
(config)#interface ce0
(config-if)#load-interval 30
(config-if)#ip address 10.0.1.26/31
(config-if)#mtu 9216
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

```
(config)#interface lo
(config-if)#ip address 30.0.1.3/32 secondary
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe4
(config-if)#speed 10g
(config-if)#load-interval 30
(config-if)#ip address 10.0.1.5/31
(config-if)#mtu 9216
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe24
(config-if)#load-interval 30
(config-if)#mtu 9600
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe24.51
(config-if)#description L3VPN-VRF51
(config-if)#encapsulation dot1q 51
(config-if)#load-interval 30
(config-if)#ip vrf forwarding VRF51
(config-if)#ip address 163.0.51.1/24
(config-if)#mtu 9216
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe25
(config-if)#load-interval 30
(config-if)#mtu 9600
(config-if)#commit
(config-if)#exit
```

```
(config)#interface xe25.6
(config-if)#encapsulation dot1q 6
(config-if)#load-interval 30
(config-if)#ip address 10.0.1.7/31
(config-if)#mtu 9216
(config-if)#label-switching
(config-if)#ip ospf network point-to-point
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

4. Configure OSPF process.

```
(config)#router ospf 100
(config-router)#ospf router-id 30.0.1.3
(config-router)#network 10.0.1.4/31 area 0.0.0.0
(config-router)#network 10.0.1.6/31 area 0.0.0.0
(config-router)#network 10.0.1.26/31 area 0.0.0.0
(config-router)#network 30.0.1.3/32 area 0.0.0.0
(config-router)#commit
(config-router)#exit
```

5. Configure the BGP.

```
(config)#router bgp 65535
(config-router)#neighbor 30.0.1.1 remote-as 65535
(config-router)#neighbor 30.0.1.1 update-source lo
(config-router)#address-family ipv4 unicast
(config-router-af)#neighbor 30.0.1.1 activate
(config-router-af)#exit-address-family
(config-router)#address-family vpnv4 unicast
(config-router-af)#neighbor 30.0.1.1 activate
(config-router-af)#exit-address-family
(config-router)#address-family ipv4 vrf VRF51
(config-router-af)#redistribute connected
(config-router-af)#exit-address-family
(config-router)#commit
(config-router)#exit
```

6. Configure RSVP trunk and enable reoptimization.

```
(config)#rsvp-trunk PE2-PE1 ipv4
(config-rsvp)#reoptimize
(config-rsvp)#to 30.0.1.1
(config-rsvp)#commit
(config-rsvp)#exit
```

Running Configuration



Note: Before configuration meet all [Prerequisites \(page 834\)](#).

P1

```
!
router rsvp
!
interface lo
 ip address 30.0.1.2/32 secondary
!
interface xe4
 load-interval 30
 ip address 10.0.1.4/31
 mtu 9216
 label-switching
 ip ospf network point-to-point
 enable-rsvp
!
interface xe44
 load-interval 30
 ip address 10.0.1.3/31
 mtu 9216
 label-switching
 ip ospf network point-to-point
 enable-rsvp
```

```
!  
router ospf 100  
  ospf router-id 30.0.1.2  
  network 10.0.1.2/31 area 0.0.0.0  
  network 10.0.1.4/31 area 0.0.0.0  
  network 30.0.1.2/32 area 0.0.0.0  
!
```

P2

```
!  
!  
router rsvp  
!  
interface ce49  
  load-interval 30  
  ip address 10.0.1.27/31  
  mtu 9216  
  label-switching  
  ip ospf network point-to-point  
  enable-rsvp  
!  
interface lo  
  ip address 30.0.1.5/32 secondary  
!  
interface xe26  
  load-interval 30  
  mtu 9600  
!  
interface xe26.20  
  encapsulation dot1q 20  
  load-interval 30  
  ip address 10.0.1.21/31  
  mtu 9216  
  label-switching  
  ip ospf network point-to-point  
  enable-rsvp  
!  
router ospf 100  
  ospf router-id 30.0.1.5  
  network 10.0.1.20/31 area 0.0.0.0  
  network 10.0.1.26/31 area 0.0.0.0  
  network 30.0.1.5/32 area 0.0.0.0  
!
```

PE1

```
!  
ip vrf VRF51  
  rd 1.1.1.1:51  
  route-target export 100:51  
  route-target import 300:51  
!  
router rsvp  
!  
interface lo  
  ip address 30.0.1.1/32 secondary  
!  
interface xe24  
  load-interval 30  
  ip address 10.0.1.2/31  
  mtu 9216
```

```
label-switching
ip ospf network point-to-point
enable-rsvp
!
interface xe25
load-interval 30
mtu 9600
!
interface xe25.6
encapsulation dot1q 6
load-interval 30
ip address 10.0.1.6/31
mtu 9216
label-switching
ip ospf network point-to-point
enable-rsvp
!
interface xe26
load-interval 30
mtu 9600
!
interface xe26.20
encapsulation dot1q 20
load-interval 30
ip address 10.0.1.20/31
mtu 9216
label-switching
ip ospf network point-to-point
enable-rsvp
!
interface xe27
load-interval 30
mtu 9600
!
interface xe27.51
description L3VPN-VRF51
encapsulation dot1q 51
load-interval 30
ip vrf forwarding VRF51
ip address 161.0.51.1/24
mtu 9216
!
exit
!
router ospf 100
ospf router-id 30.0.1.1
network 10.0.1.2/31 area 0.0.0.0
network 10.0.1.6/31 area 0.0.0.0
network 10.0.1.20/31 area 0.0.0.0
network 30.0.1.1/32 area 0.0.0.0
!
router bgp 65535
neighbor 30.0.1.3 remote-as 65535
neighbor 30.0.1.3 update-source lo
!
address-family ipv4 unicast
neighbor 30.0.1.3 activate
exit-address-family
!
address-family vpnv4 unicast
neighbor 30.0.1.3 activate
exit-address-family
!
address-family ipv4 vrf VRF51
redistribute connected
exit-address-family
!
exit
```

```
!  
rsvp-trunk PE1-PE2 ipv4  
  reoptimize  
  to 30.0.1.3  
!
```

PE2

```
!  
ip vrf VRF51  
  rd 3.3.3.3:51  
  route-target import 100:51  
  route-target export 300:51  
!  
router rsvp  
!  
interface ce0  
  load-interval 30  
  ip address 10.0.1.26/31  
  mtu 9216  
  label-switching  
  ip ospf network point-to-point  
  enable-rsvp  
!  
interface lo  
  ip address 30.0.1.3/32 secondary  
!  
interface xe4  
  speed 10g  
  load-interval 30  
  ip address 10.0.1.5/31  
  mtu 9216  
  label-switching  
  ip ospf network point-to-point  
  enable-rsvp  
!  
interface xe24  
  load-interval 30  
  mtu 9600  
!  
interface xe24.51  
  description L3VPN-VRF51  
  encapsulation dot1q 51  
  load-interval 30  
  ip vrf forwarding VRF51  
  ip address 163.0.51.1/24  
  mtu 9216  
!  
interface xe25  
  load-interval 30  
  mtu 9600  
!  
interface xe25.6  
  encapsulation dot1q 6  
  load-interval 30  
  ip address 10.0.1.7/31  
  mtu 9216  
  label-switching  
  ip ospf network point-to-point  
  enable-rsvp  
!  
router ospf 100  
  ospf router-id 30.0.1.3  
  network 10.0.1.4/31 area 0.0.0.0  
  network 10.0.1.6/31 area 0.0.0.0
```

```

network 10.0.1.26/31 area 0.0.0.0
network 30.0.1.3/32 area 0.0.0.0
!
router bgp 65535
neighbor 30.0.1.1 remote-as 65535
neighbor 30.0.1.1 update-source lo
!
address-family ipv4 unicast
neighbor 30.0.1.1 activate
exit-address-family
!
address-family vpnv4 unicast
neighbor 30.0.1.1 activate
exit-address-family
!
address-family ipv4 vrf VRF51
redistribute connected
exit-address-family
!
exit
!
rsvp-trunk PE2-PE1 ipv4
reoptimize
to 30.0.1.1
!

```

Validation

Base Configuration Validations

- Verify OSPF adjacency is up on all nodes.
- Verify BGP peering , RSVP session and L3VPN service are up between PE1 & PE2.

PE1

```
PE1#show ip ospf neighbor
```

Total number of full neighbors: 3

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|-----------|-----------|-------------|
| 30.0.1.2 | 1 | Full/ - | 00:00:38 | 10.0.1.3 | xe24 | 0 |
| 30.0.1.3 | 1 | Full/ - | 00:00:37 | 10.0.1.7 | xe25.6 | 0 |
| 30.0.1.5 | 1 | Full/ - | 00:00:38 | 10.0.1.21 | xe26.20 | 0 |

PE1#

```
PE1#show ip bgp summary
```

BGP router identifier 30.0.1.1, local AS number 65535

BGP table version is 2

1 BGP AS-PATH entries

0 BGP community entries

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|-------|--------|--------|--------|-----|------|----------|--------------|------|
| 30.0.1.3 | 4 | 65535 | 13 | 14 | 2 | 0 | 0 | 00:04:26 | 0 | |

Total number of neighbors 1

Total number of Established sessions 1

PE1#

```
PE1#show rsvp session
```

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass

State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary

* indicates the session is active with local repair at one or more nodes

(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|----------|----------|----------|---------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.3 | 30.0.1.1 | 5001 | 2201 | PRI | PE1-PE2- | |
| Primary | UP | 00:00:15 | 1 1 SE | - | 24320 | |

Total 1 displayed, Up 1, Down 0.

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|----------|----------|----------|---------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.1 | 30.0.1.3 | 5001 | 2201 | PRI | PE2-PE1- | |
| Primary | UP | 00:00:06 | 1 1 SE | 24320 | - | |

Total 1 displayed, Up 1, Down 0.

PE1#

PE1#show mpls vrf-forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN

(m) - Service mapped over multipath transport

(e) - Service mapped over ECMP

(D) - Down

Ext-Color - Extended-community color advertised by BGP

B(x) - BGP EVPN MPLS Services

| Code | FEC | Nexthop | UpTime | FTN-ID | VRF-ID | Nhlfe-ID | Pri | Out- |
|-------|---------------|----------|--------|-----------|--------|----------|-----|-------|
| Label | Out-Intf | | | Ext-Color | | | | |
| B> | 163.0.51.0/24 | | | 1 | 2 | 4 | - | - |
| | - | 00:00:30 | - | | | 2 | Yes | 24960 |
| | 30.0.1.3 | - | - | | | | | - |

PE1#

P1

P1#show ip ospf neighbor

Total number of full neighbors: 2

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|----------|-----------|-------------|
| 30.0.1.1 | 1 | Full/ - | 00:00:35 | 10.0.1.2 | xe44 | 0 |
| 30.0.1.3 | 1 | Full/ - | 00:00:33 | 10.0.1.5 | xe4 | 0 |

P1#

P2

P2#show ip ospf neighbor

Total number of full neighbors: 0

OSPF process 0 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------|-----------|---------|-----------|-------------|
|-------------|-----|-------|-----------|---------|-----------|-------------|

Total number of full neighbors: 2

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|-----------|-----------|-------------|
| 30.0.1.1 | 1 | Full/ - | 00:00:39 | 10.0.1.20 | xe26.20 | 0 |
| 30.0.1.3 | 1 | Full/ - | 00:00:36 | 10.0.1.26 | ce49 | 0 |

PE2

PE2#show ip ospf neighbor

Total number of full neighbors: 3

OSPF process 100 VRF(default):

```

Neighbor ID      Pri   State           Dead Time   Address      Interface     Instance ID
30.0.1.2         1    Full/ -        00:00:32    10.0.1.4     xe4           0
30.0.1.1         1    Full/ -        00:00:29    10.0.1.6     xe25.6        0
30.0.1.5         1    Full/ -        00:00:30    10.0.1.27    ce0           0
PE2#
PE2#show ip bgp summary
BGP router identifier 10.0.1.26, local AS number 65535
BGP table version is 2
1 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS      MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd   Desc
30.0.1.1          4  65535      13       13       2       0     0    00:04:33         0

Total number of neighbors 1

Total number of Established sessions 1

PE2#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime    Rt  Style  Labelin  Labelout
30.0.1.1   30.0.1.3   5001     2201     PRI    PE2-PE1-         24321
Primary              UP    00:00:10  1 1 SE    -
Total 1 displayed, Up 1, Down 0.

Egress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime    Rt  Style  Labelin  Labelout
30.0.1.3   30.0.1.1   5001     2201     PRI    PE1-PE2-         24320
Primary              UP    00:00:20  1 1 SE    -
Total 1 displayed, Up 1, Down 0.

PE2#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code   FEC
Label  Out-Intf      Nexthop      UpTime      FTN-ID VRF-ID   Nhlfe-ID   Pri   Out-
B>    161.0.51.0/24
      -              00:00:15    -              1       2         4         -    -
      30.0.1.1      -              -              -              -       -         2         Yes  24960
PE2#

```

Scenario 1: Verify Reoptimization for Primary Tunnels

1. Configure re-optimization timers under global RSVPmode.

By default LSP re-optimization timer is 5 minutes.

```

PE1(config)#router rsvp
PE1(config-router)#lsp-reoptimization-timer 4

```

2. Shutdown the best path between PE1 and PE2.

```

PE1(config)#int xe25
PE1(config-if)#shut
PE1(config-if)#commit

```

3. Enable reoptimization under trunk.

```
PE1(config)#rsvp-trunk PE1-PE2
PE1(config-trunk)#reoptimize
```

```
PE1#show running-config rsvp
!
router rsvp
  lsp-reoptimization-timer 4
!
interface xe24
  enable-rsvp
!
interface xe25.6
  enable-rsvp
!
interface xe26.20
  enable-rsvp
!
rsvp-trunk PE1-PE2 ipv4
  reoptimize
  to 30.0.1.3
!
!
!
```

```
PE1#show cspf lsp
  LSP Name      : PE1-PE2-Primary
  Trunk ID      : 5001
  Lsp Id        : 2201
  OSPF ID       : 100
  Ingress       : 30.0.1.1
  Egress        : 30.0.1.3
  Ext Tunnel ID : 30.0.1.1
  LSP Type      : 0
  Client ID     : 1
  State         : Established(3)
  Setup Priority : 7
  Hold Priority  : 0
  Hop Limit     : 255
  Include Mask   : 0x0
  Exclude Mask   : 0x0
  Include All Mask : 0x0
  LSP Metric     : 2
  LSP Reoptimize : Yes
  LSP in ReopList : Yes
  LSP ReopTime   : 4 minutes (Offset 17 seconds)
  Next Reop Due in: 3 minutes 50 seconds
  Computed ERO   :
    10.0.1.21
    10.0.1.26
```

```
PE1#
```

```
PE2(config)#rsvp-trunk PE2-PE1
PE2(config-trunk)#reoptimize
```

```
PE2#show running-config rsvp
!
router rsvp
!
interface ce0
  enable-rsvp
!
interface xe4
  enable-rsvp
!
```

```

interface xe25.6
  enable-rsvp
  !
  rsvp-trunk PE2-PE1 ipv4
  reoptimize
  to 30.0.1.1
  !

```

```

PE2#show cspf lsp
  LSP Name       : PE2-PE1-Primary
  Trunk ID       : 5001
  Lsp Id         : 2201
  OSPF ID        : 100
  Ingress        : 30.0.1.3
  Egress         : 30.0.1.1
  Ext Tunnel ID  : 30.0.1.3
  LSP Type       : 0
  Client ID      : 1
  State          : Established(3)
  Setup Priority  : 7
  Hold Priority   : 0
  Hop Limit      : 255
  Include Mask   : 0x0
  Exclude Mask   : 0x0
  Include All Mask : 0x0
  LSP Metric     : 2
  LSP Reoptimize : Yes
  LSP in ReopList : Yes
  LSP ReopTime   : 5 minutes (Offset 17 seconds)
  Next Reop Due in: 4 minutes 59 seconds
  Computed ERO   :
    10.0.1.27
    10.0.1.20

```

PE2#

4. Verify traffic counters between PE1 and PE2 before reoptimizing the tunnel. Verify that traffic is not going via the directly connected link between PE1-PE2.

```
PE1#show interface counters rate mbps
```

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-----------|---------|--------|---------|--------|
| xe26 | 498.80 | 13343 | 499.40 | 13342 |
| xe27 | 498.55 | 13342 | 497.95 | 13343 |

PE1#

```
P2#show interface counters rate mbps
```

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-----------|---------|--------|---------|--------|
| ce49 | 498.73 | 13343 | 498.25 | 13343 |
| xe26 | 498.68 | 13343 | 499.16 | 13343 |

P2#

```
PE2#show interface counters rate mbps
```

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-----------|---------|--------|---------|--------|
| ce0 | 488.31 | 13069 | 488.72 | 13068 |
| xe24 | 488.30 | 13068 | 487.90 | 13068 |

PE2#

5. Bring up the directly connected path between PE1 and PE2 & wait for re-optimization timer expiry.

```

PE1(config)#int xe25
PE1(config-if)#no shut
PE1(config-if)#commit

```

```
PE1#show ip ospf neighbor
```

```

Total number of full neighbors: 3
OSPF process 100 VRF(default):

```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|-----------|-----------|-------------|
| 30.0.1.2 | 1 | Full/ - | 00:00:31 | 10.0.1.3 | xe24 | 0 |
| 30.0.1.3 | 1 | Full/ - | 00:00:34 | 10.0.1.7 | xe25.6 | 0 |
| 30.0.1.5 | 1 | Full/ - | 00:00:28 | 10.0.1.21 | xe26.20 | 0 |

```
PE1#
```

```
PE2#show ip ospf neighbor
```

```

Total number of full neighbors: 3
OSPF process 100 VRF(default):

```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|-----------|-----------|-------------|
| 30.0.1.2 | 1 | Full/ - | 00:00:31 | 10.0.1.4 | xe4 | 0 |
| 30.0.1.1 | 1 | Full/ - | 00:00:38 | 10.0.1.6 | xe25.6 | 0 |
| 30.0.1.5 | 1 | Full/ - | 00:00:31 | 10.0.1.27 | ce0 | 0 |

```
PE2#
```

```
PE2#show cspf lsp
```

```

LSP Name       : PE2-PE1-Primary
Trunk ID       : 5001
Lsp Id         : 2201
OSPF ID        : 100
Ingress        : 30.0.1.3
Egress         : 30.0.1.1
Ext Tunnel ID  : 30.0.1.3
LSP Type       : 0
Client ID      : 1
State          : Established(3)
Setup Priority  : 7
Hold Priority   : 0
Hop Limit      : 255
Include Mask    : 0x0
Exclude Mask    : 0x0
Include All Mask : 0x0
LSP Metric     : 2
LSP Reoptimize : Yes
LSP in ReopList : Yes
LSP ReopTime   : 5 minutes (Offset 17 seconds)
Next Reop Due in: 24 seconds
Computed ERO    :
                  10.0.1.27
                  10.0.1.20

```

```
PE2#
```

After re-optimization timer expiry , tunnel should come up via the best path.

```

PE2#2025 Oct 07 11:43:33.825 : PE2 : RSVP : NOTIF : [RSVP_OPR_MBB_SESSION_4]: MBB session being
created for RSVP Session PE2-PE1-Primary ingress:30.0.1.3/32 egress:30.0.1.1/32
2025 Oct 07 11:43:33.829 : PE2 : RSVP : NOTIF : [RSVP_OPR_MBB_SESSION_4]: MBB session PE2-PE1-Primary
with LSP ID 2202 has come up, delete the old session with LSP ID 2201
2025 Oct 07 11:43:33.830 : PE2 : RSVP : NOTIF : [RSVP_OPR_SESSION_STATUS_DOWN_4]: Ingress RSVP
session PE2-PE1-Primary ingress:30.0.1.3/32 egress:30.0.1.1/32 tunnel id 5001 lsp id 2201 changed
state to down
2025 Oct 07 11:43:33.830 : PE2 : RSVP : NOTIF : [RSVP_OPR_SESSION_STATUS_UP_4]: Ingress RSVP session
PE2-PE1-Primary ingress:30.0.1.3/32 egress:30.0.1.1/32 tunnel id 5001 lsp id 2202 changed state to up

```

```

PE2#show cspf lsp
  LSP Name       : PE2-PE1-Primary
  Trunk ID       : 5001
  Lsp Id         : 2202
  OSPF ID        : 100
  Ingress        : 30.0.1.3
  Egress         : 30.0.1.1
  Ext Tunnel ID  : 30.0.1.3
  LSP Type       : 0
  Client ID      : 1
  State          : Established(3)
  Setup Priority  : 7
  Hold Priority   : 0
  Hop Limit      : 255
  Include Mask   : 0x0
  Exclude Mask   : 0x0
  Include All Mask : 0x0
  LSP Metric     : 1
  LSP Reoptimize : Yes
  LSP in ReopList : Yes
  LSP ReopTime   : 5 minutes (Offset 18 seconds)
  Next Reop Due in: 5 minutes 15 seconds
  Computed ERO   :
                  10.0.1.6

PE2#

```

6. Verify traffic is up on the best path after re-optimization.

```

PE1#show interface counters rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
| xe25      | 499.66  | 13343  | 499.67  | 13343  |
| xe27      | 498.81  | 13343  | 498.81  | 13343  |
PE1#

```

```

PE2#show interface counters rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
| xe24      | 498.63  | 13344  | 497.74  | 13346  |
| xe25      | 498.59  | 13346  | 499.48  | 13345  |
PE2#

```

Scenario 2: Verify Re-optimization for Secondary Tunnels

1. Configure secondary on RSVP trunks.

```

PE1(config)#rsvp-trunk PE1-PE2
PE1(config-trunk)#secondary
PE1(config-trunk)#commit

```

```

PE1#show run rsvp
!
router rsvp
  lsp-reoptimization-timer 4
!
!
interface xe24
  enable-rsvp
!
interface xe25.6
  enable-rsvp
!

```

```

interface xe26.20
  enable-rsvp
  !
  !
  !
  !
  !
  rsvp-trunk PE1-PE2 ipv4
    reoptimize
    secondary
    to 30.0.1.3
  !
  !

```

```

PE2(config)#rsvp-trunk PE2-PE1
PE2(config-trunk)#secondary
PE2(config-trunk)#commit

```

```

PE2#show run rsvp
!
router rsvp
!
!
interface ce0
  enable-rsvp
!
interface xe4
  enable-rsvp
!
interface xe25.6
  enable-rsvp
!
!
!
!
!
rsvp-trunk PE2-PE1 ipv4
  reoptimize
  secondary
  to 30.0.1.1
!
!
!
PE2#

```

2. Verify secondary tunnel paths.

```

PE1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|----------------------------------|----------|--------|----------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.3 | 30.0.1.1 | 5001 | 2201 | SEC | PE1-PE2- | |
| Secondary | | UP | 00:00:18 | 1 1 SE | - | 24320 |
| 30.0.1.3 | 30.0.1.1 | 5001 | 2202 | PRI | PE1-PE2- | |
| Primary | | UP | 00:03:30 | 1 1 SE | - | 24321 |
| Total 2 displayed, Up 2, Down 0. | | | | | | |

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-----------|----------|--------|----------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.1 | 30.0.1.3 | 5001 | 2201 | PRI | PE2-PE1- | |
| Secondary | | UP | 00:00:09 | 1 1 SE | 24320 | - |

```

30.0.1.1      30.0.1.3      5001      2202      PRI      PE2-PE1-
Primary      UP      00:02:16  1 1 SE      24321      -
Total 2 displayed, Up 2, Down 0.

```

```

PE1#show cspf lsp
  LSP Name      : PE1-PE2-Primary
  Trunk ID      : 5001
  Lsp Id        : 2202
  OSPF ID       : 100
  Ingress       : 30.0.1.1
  Egress        : 30.0.1.3
  Ext Tunnel ID : 30.0.1.1
  LSP Type      : 0
  Client ID     : 1
  State         : Established(3)
  Setup Priority : 7
  Hold Priority  : 0
  Hop Limit     : 255
  Include Mask  : 0x0
  Exclude Mask  : 0x0
  Include All Mask: 0x0
  LSP Metric    : 1
  LSP Reoptimize : Yes
  LSP in ReopList : Yes
  LSP ReopTime  : 4 minutes (Offset 18 seconds)
  Next Reop Due in: 29 seconds
  Computed ERO  :
    10.0.1.7

  LSP Name      : PE1-PE2-Secondary
  Trunk ID      : 5001
  Lsp Id        : 2201
  OSPF ID       : 100
  Ingress       : 30.0.1.1
  Egress        : 30.0.1.3
  Ext Tunnel ID : 30.0.1.1
  LSP Type      : 0
  Client ID     : 1
  State         : Established(3)
  Setup Priority : 7
  Hold Priority  : 0
  Hop Limit     : 255
  Include Mask  : 0x0
  Exclude Mask  : 0x0
  Include All Mask: 0x0
  LSP Metric    : 2
  LSP Reoptimize : Yes
  LSP in ReopList : Yes
  LSP ReopTime  : 4 minutes (Offset 17 seconds)
  Next Reop Due in: 3 minutes 40 seconds
  Exclude Path Constraint :
    10.0.1.7      exclude link
  Computed ERO  :
    10.0.1.21
    10.0.1.26

```

```
PE1#
```

3. Modify the te-metric on secondary tunnel path and reoptimize the tunnel.

```

PE1(config)#interface xe26.20
PE1(config-if)#te-metric 1000
PE1(config-if)#commit
PE1(config-if)#end

```

```
PE1#show run interface xe26.20
```



```

!
interface xe26.20
 encapsulation dot1q 20
 load-interval 30
 ip address 10.0.1.20/31
 mtu 9216
 label-switching
 ip ospf network point-to-point
 te-metric 1000
 enable-rsvp
!
!
PE1#

```

```

PE1#show cspf lsp PE1-PE2-Secondary
  LSP Name       : PE1-PE2-Secondary
  Trunk ID       : 5001
  Lsp Id         : 2201
  OSPF ID        : 100
  Ingress        : 30.0.1.1
  Egress         : 30.0.1.3
  Ext Tunnel ID  : 30.0.1.1
  LSP Type       : 0
  Client ID      : 1
  State          : Established(3)
  Setup Priority  : 7
  Hold Priority   : 0
  Hop Limit      : 255
  Include Mask   : 0x0
  Exclude Mask   : 0x0
  Include All Mask : 0x0
  LSP Metric     : 2
  LSP Reoptimize : Yes
  LSP in ReopList : Yes
  LSP ReopTime   : 4 minutes (Offset 17 seconds)
  Next Reop Due in: 1 seconds
  Exclude Path Constraint :
    10.0.1.7      exclude link
  Computed ERO   :
    10.0.1.21
    10.0.1.26

```

```

PE1#2025 Oct 07 11:59:02.192 : PE1 : RSVP : NOTIF : [RSVP_OPR_MBB_SESSION_4]: MBB session being
created for RSVP Session PE1-PE2-Secondary ingress:30.0.1.1/32 egress:30.0.1.3/32
2025 Oct 07 11:59:02.203 : PE1 : RSVP : NOTIF : [RSVP_OPR_MBB_SESSION_4]: MBB session PE1-PE2-
Secondary with LSP ID 2203 has come up, delete the old session with LSP ID 2201
2025 Oct 07 11:59:02.203 : PE1 : RSVP : NOTIF : [RSVP_OPR_SESSION_STATUS_DOWN_4]: Ingress RSVP
session PE1-PE2-Secondary ingress:30.0.1.1/32 egress:30.0.1.3/32 tunnel id 5001 lsp id 2201 changed
state to down
2025 Oct 07 11:59:02.203 : PE1 : RSVP : NOTIF : [RSVP_OPR_SESSION_STATUS_UP_4]: Ingress RSVP session
PE1-PE2-Secondary ingress:30.0.1.1/32 egress:30.0.1.3/32 tunnel id 5001 lsp id 2203 changed state to
up
PE1#

```

4. Validate secondary tunnel path after re-optimization.

```

PE1#show cspf lsp PE1-PE2-Secondary
  LSP Name       : PE1-PE2-Secondary
  Trunk ID       : 5001
  Lsp Id         : 2203
  OSPF ID        : 100
  Ingress        : 30.0.1.1
  Egress         : 30.0.1.3
  Ext Tunnel ID  : 30.0.1.1
  LSP Type       : 0

```

```

Client ID       : 1
State           : Established(3)
Setup Priority   : 7
Hold Priority    : 0
Hop Limit       : 255
Include Mask    : 0x0
Exclude Mask    : 0x0
Include All Mask: 0x0
LSP Metric      : 2
LSP Reoptimize  : Yes
LSP in ReopList : Yes
LSP ReopTime    : 4 minutes (Offset 19 seconds)
Next Reop Due in: 4 minutes 16 seconds
Exclude Path Constraint :
    10.0.1.7    exclude link
Computed ERO    :
    10.0.1.3
    10.0.1.5

```

PE1#

Scenario 3: Verify Reoptimization for Detour Tunnels



Note: Reoptimization is implicitly enabled for detour tunnels. Detour tunnel does not reoptimize when it is in use.

1. Enable one-one protection under trunk.

```

PE1(config)#rsvp-trunk PE1-PE2
PE1(config-trunk)#primary fast-reroute protection one-to-one
PE1(config-trunk)#primary fast-reroute node-protection
PE1(config-trunk)#commit

```

```

PE1#show run rsvp
!
router rsvp
  lsp-reoptimization-timer 4
!
interface xe24
  enable-rsvp
!
interface xe25.6
  enable-rsvp
!
interface xe26.20
  enable-rsvp
!
rsvp-trunk PE1-PE2 ipv4
  primary fast-reroute protection one-to-one
  primary fast-reroute node-protection
  to 30.0.1.3
!

```

2. Validate current path taken by detour tunnel .

```

PE1#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

Ingress RSVP:

| To | Uptime | Rt | From | Style | Labelin | Labelout | Tun-ID | LSP-ID | Type | LSPName | State |
|----------|--------|----|-----------|-------|---------|----------|--------|--------|------|----------|-------|
| 30.0.1.3 | | | 30.0.1.1 | | | | 5001 | 2203 | PRI | PE1-PE2- | |
| Primary | | | | UP | | 00:00:20 | 1 1 SE | | - | 24320 | |
| 30.0.1.3 | | | 10.0.1.20 | | | | 5001 | 2203 | DTR | PE1-PE2- | |
| Detour | | | | UP | | 00:00:20 | 1 1 SE | | - | 24320 | |

Total 2 displayed, Up 2, Down 0.

Egress RSVP:

| To | Uptime | Rt | From | Style | Labelin | Labelout | Tun-ID | LSP-ID | Type | LSPName | State |
|----------|--------|----|-----------|-------|---------|----------|--------|--------|-------|----------|-------|
| 30.0.1.1 | | | 30.0.1.3 | | | | 5001 | 2203 | PRI | PE2-PE1- | |
| Primary | | | | UP | | 00:00:08 | 1 1 SE | | 24320 | - | |
| 30.0.1.1 | | | 10.0.1.26 | | | | 5001 | 2203 | PRI | PE2-PE1- | |
| Detour | | | | UP | | 00:00:08 | 1 1 SE | | 24321 | - | |

Total 2 displayed, Up 2, Down 0.

PE1#show cspf lsp

```

LSP Name       : PE1-PE2-Primary
Trunk ID       : 5001
Lsp Id         : 2203
OSPF ID        : 100
Ingress        : 30.0.1.1
Egress         : 30.0.1.3
Ext Tunnel ID  : 30.0.1.1
LSP Type       : 0
Client ID      : 1
State          : Established(3)
Setup Priority  : 7
Hold Priority   : 0
Hop Limit      : 255
Include Mask   : 0x0
Exclude Mask   : 0x0
Include All Mask : 0x0
LSP Metric     : 1
LSP Reoptimize : No
LSP in ReopList : No
Computed ERO   :
    10.0.1.7

LSP Name       : PE1-PE2-Detour
Trunk ID       : 5001
Lsp Id         : 2203
OSPF ID        : 100
Ingress        : 10.0.1.20
Egress         : 30.0.1.3
Ext Tunnel ID  : 30.0.1.1
LSP Type       : 0
Client ID      : 1
State          : Established(3)
Setup Priority  : 7
Hold Priority   : 0
Hop Limit      : 255
Include Mask   : 0x0
Exclude Mask   : 0x0
Include All Mask : 0x0
LSP Metric     : 2
LSP Reoptimize : Yes
LSP in ReopList : Yes
LSP ReopTime   : 4 minutes (Offset 19 seconds)
Next Reop Due in: 3 minutes 52 seconds
Exclude Path Constraint :
    10.0.1.7      exclude link
Computed ERO   :
    10.0.1.21
    10.0.1.26

```

```
PE1#
```

3. Reoptimize detour tunnel by modifying te-metric.

```
PE1(config)#int xe26.20
PE1(config-if)#te-metric 10000
PE1(config-if)#commit
```

```
PE1#show run interface xe26.20
!
interface xe26.20
 encapsulation dot1q 20
 load-interval 30
 ip address 10.0.1.20/31
 mtu 9216
 label-switching
 ip ospf network point-to-point
 te-metric 10000
 enable-rsvp
!
!
```

```
PE1#show cspf lsp PE1-PE2-Detour
LSP Name       : PE1-PE2-Detour
Trunk ID       : 5001
Lsp Id         : 2203
OSPF ID        : 100
Ingress        : 10.0.1.20
Egress         : 30.0.1.3
Ext Tunnel ID  : 30.0.1.1
LSP Type       : 0
Client ID      : 1
State          : Established(3)
Setup Priority  : 7
Hold Priority   : 0
Hop Limit      : 255
Include Mask   : 0x0
Exclude Mask   : 0x0
Include All Mask : 0x0
LSP Metric     : 2
LSP Reoptimize : Yes
LSP in ReopList : Yes
LSP ReopTime   : 4 minutes (Offset 19 seconds)
Next Reop Due in: 1 minutes 34 seconds
Exclude Path Constraint :
  10.0.1.7      exclude link
Computed ERO   :
  10.0.1.21
  10.0.1.26
```

```
PE1#
```

```
PE1#rsvp-session PE1-PE2-Detour force-reoptimize
PE1#2025 Oct 07 12:04:09.161 : PE1 : RSVP : CRITI : [RSVP_OPR_SESSION_STATUS_DOWN_2]: Ingress RSVP
session PE1-PE2-Detour ingress:10.0.1.20/32 egress:30.0.1.3/32 tunnel id 5001 lsp id 2203 changed
state to down
2025 Oct 07 12:04:09.172 : PE1 : RSVP : NOTIF : [RSVP_OPR_SESSION_STATUS_UP_4]: Ingress RSVP session
PE1-PE2-Detour ingress:10.0.1.2/32 egress:30.0.1.3/32 tunnel id 5001 lsp id 2203 changed state to up
```

```
PE1#
```

4. Validate detour tunnel path after reoptimization.

```
PE1#show cspf lsp PE1-PE2-Detour
```

```

LSP Name       : PE1-PE2-Detour
Trunk ID       : 5001
Lsp Id        : 2203
OSPF ID       : 100
Ingress       : 10.0.1.2
Egress        : 30.0.1.3
Ext Tunnel ID  : 30.0.1.1
LSP Type      : 0
Client ID     : 1
State         : Established(3)
Setup Priority : 7
Hold Priority  : 0
Hop Limit     : 255
Include Mask   : 0x0
Exclude Mask   : 0x0
Include All Mask : 0x0
LSP Metric    : 2
LSP Reoptimize : Yes
LSP in ReopList : Yes
LSP ReopTime   : 4 minutes (Offset 19 seconds)
Next Reop Due in: 4 minutes 9 seconds
Exclude Path Constraint :
    10.0.1.7    exclude link
Computed ERO   :
    10.0.1.3
    10.0.1.5

```

Scenario 4: Verify Reoptimization for Auto-bypass Tunnels



Note:

1. Auto-bypass tunnels do not reoptimize when they are in use. This behavior is same with manual bypass tunnels too.
2. Reoptimization for auto-bypass tunnels are enabled under global auto-bypass configuration in router rsvp where as reoptimization for manual bypass tunnels are enabled under manual bypass tunnel configuration mode.

1. Enable reoptimization for auto-bypass and enable facility bypass protection under rsvp trunk.

```

PE1(config)#router rsvp
PE1(config-router)#auto-bypass
PE1(config-auto-bypass)#enable
PE1(config-auto-bypass)#inactivity-timer 3
PE1(config-auto-bypass)#attributes best-effort
PE1(config-auto-bypass-attr)#protection-capability node
PE1(config-auto-bypass-attr)#commit
PE1(config-auto-bypass-attr)#end

PE1(config)#router rsvp
PE1(config-router)#auto-bypass
PE1(config-auto-bypass)#attributes best-effort
PE1(config-auto-bypass-attr)#reoptimize
PE1(config-auto-bypass-attr)#commit
PE1(config-auto-bypass-attr)#end

PE1(config)#rsvp-trunk PE1-PE2
PE1(config-trunk)#primary fast-reroute protection facility
PE1(config-trunk)#primary fast-reroute node-protection
PE1(config-trunk)#reoptimize
PE1(config-trunk)#commit

```

```

PE1#show run rsvp
!
router rsvp
 lsp-reoptimization-timer 4
 auto-bypass
  attributes best-effort
  protection-capability node
  reoptimize
  exit
 inactivity-timer 3
 enable
 exit
!
interface xe24
 enable-rsvp
!
interface xe25.6
 enable-rsvp
!
interface xe26.20
 enable-rsvp
!
rsvp-trunk PE1-PE2 ipv4
 reoptimize
 primary fast-reroute protection facility
 primary fast-reroute node-protection
 to 30.0.1.3
!

```

2. Validate path taken by auto-bypass tunnel.

```

PE1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|----------|----------|--------|----------|----------|---------------------------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.3 | 30.0.1.1 | 5001 | 2204 | PRI | PE1-PE2- | |
| Primary | | UP | 00:00:34 | 1 1 SE | - | 24320 |
| 30.0.1.3 | 30.0.1.1 | 5002 | 2202 | BPS | BL-10.0.1.7-30.0.1.3-100-Bypass | |
| UP | 00:00:34 | 1 1 SE | - | | 24320 | |

Total 2 displayed, Up 2, Down 0.

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|----------|----------|--------|----------|----------|---------------------------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.1 | 30.0.1.3 | 5001 | 2204 | PRI | PE2-PE1- | |
| Primary | | UP | 00:00:16 | 1 1 SE | 24320 | - |
| 30.0.1.1 | 30.0.1.3 | 5002 | 2202 | PRI | BL-10.0.1.6-30.0.1.1-100-Bypass | |
| UP | 00:00:16 | 1 1 SE | 3 | | - | |

Total 2 displayed, Up 2, Down 0.

PE1#

```

PE1#show rsvp bypass protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
              3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

```

```

Bypass trunk: BL-10.0.1.7-30.0.1.3-100
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:

```

| Tunnel-id | Lsp-Id | Lsp-Name | Role | Ext_tnl_ |
|-----------|---------|----------|------------|-----------|
| id | Ingress | Egress | Match-Code | Mapped-BW |

```

5001      2204      PE1-PE2-
Primary      Ingress      30.0.1.1      30.0.1.1      30.0.1.3      0      NA
Total LSP protected : 1
Total LSPs using protection path : 0
Bandwidth in use : 0

PE1#

```

```

PE1#show cspf lsp
  LSP Name      : PE1-PE2-Primary
  Trunk ID      : 5001
  Lsp Id        : 2204
  OSPF ID       : 100
  Ingress       : 30.0.1.1
  Egress        : 30.0.1.3
  Ext Tunnel ID : 30.0.1.1
  LSP Type      : 0
  Client ID     : 1
  State         : Established(3)
  Setup Priority : 7
  Hold Priority  : 0
  Hop Limit     : 255
  Include Mask  : 0x0
  Exclude Mask  : 0x0
  Include All Mask: 0x0
  LSP Metric    : 1
  LSP Reoptimize : Yes
  LSP in ReopList : Yes
  LSP ReopTime  : 4 minutes
  Next Reop Due in: 2 minutes 31 seconds
  Computed ERO  :
    10.0.1.7

  LSP Name      : BL-10.0.1.7-30.0.1.3-100-Bypass
  Trunk ID      : 5002
  Lsp Id        : 2202
  OSPF ID       : 100
  Ingress       : 30.0.1.1
  Egress        : 30.0.1.3
  Ext Tunnel ID : 30.0.1.1
  LSP Type      : 2
  Client ID     : 1
  State         : Established(3)
  Setup Priority : 7
  Hold Priority  : 0
  Hop Limit     : 255
  Include Mask  : 0x0
  Exclude Mask  : 0x0
  Include All Mask: 0x0
  LSP Metric    : 2
  LSP Reoptimize : Yes
  LSP in ReopList : Yes
  LSP ReopTime  : 4 minutes (Offset 14 seconds)
  Next Reop Due in: 2 minutes 45 seconds
  Exclude Path Constraint :
    10.0.1.7      exclude link
  Computed ERO  :
    10.0.1.21
    10.0.1.26

```

3. Reoptimize auto-bypass tunnel by modifying te-metric.

```

P2(config)#interface ce49
P2(config-if)#te-metric 65535
P2(config-if)#commit
P2(config-if)#end

```

```
P2#show run interface ce49
!
interface ce49
 load-interval 30
 ip address 10.0.1.27/31
 mtu 9216
 label-switching
 ip ospf network point-to-point
 te-metric 65535
 enable-rsvp
!
```

```
PE1#show cspf lsp BL-10.0.1.7-30.0.1.3-100-Bypass
LSP Name       : BL-10.0.1.7-30.0.1.3-100-Bypass
Trunk ID       : 5002
Lsp Id         : 2202
OSPF ID        : 100
Ingress        : 30.0.1.1
Egress         : 30.0.1.3
Ext Tunnel ID  : 30.0.1.1
LSP Type       : 2
Client ID      : 1
State          : Established(3)
Setup Priority  : 7
Hold Priority   : 0
Hop Limit      : 255
Include Mask   : 0x0
Exclude Mask   : 0x0
Include All Mask : 0x0
LSP Metric     : 2
LSP Reoptimize : Yes
LSP in ReopList : Yes
LSP ReopTime   : 4 minutes (Offset 14 seconds)
Next Reop Due in: 2 seconds
Exclude Path Constraint :
  10.0.1.7    exclude link
Computed ERO   :
  10.0.1.21
  10.0.1.26
```

PE1#

```
PE1#2025 Oct 07 12:12:53.461 : PE1 : RSVP : NOTIF : [RSVP_OPR_MBB_SESSION_4]: MBB session BL-
10.0.1.7-30.0.1.3-100-Bypass with LSP ID 2203 has come up, delete the old session with LSP ID 2202
2025 Oct 07 12:12:53.461 : PE1 : RSVP : NOTIF : [RSVP_OPR_SESSION_STATUS_DOWN_4]: Ingress RSVP
session BL-10.0.1.7-30.0.1.3-100-Bypass ingress:30.0.1.1/32 egress:30.0.1.3/32 tunnel id 5002 lsp id
2202 changed state to down
2025 Oct 07 12:12:53.462 : PE1 : RSVP : NOTIF : [RSVP_OPR_SESSION_STATUS_UP_4]: Ingress RSVP session
BL-10.0.1.7-30.0.1.3-100-Bypass ingress:30.0.1.1/32 egress:30.0.1.3/32 tunnel id 5002 lsp id 2203
changed state to up
```

PE1#

4. Validate path taken by auto-bypass tunnel after reoptimization.

```
PE1#show cspf lsp BL-10.0.1.7-30.0.1.3-100-Bypass
LSP Name       : BL-10.0.1.7-30.0.1.3-100-Bypass
Trunk ID       : 5002
Lsp Id         : 2203
OSPF ID        : 100
Ingress        : 30.0.1.1
Egress         : 30.0.1.3
Ext Tunnel ID  : 30.0.1.1
LSP Type       : 2
Client ID      : 1
State          : Established(3)
```



```

Setup Priority   : 7
Hold Priority    : 0
Hop Limit       : 255
Include Mask    : 0x0
Exclude Mask    : 0x0
Include All Mask: 0x0
LSP Metric      : 2
LSP Reoptimize  : Yes
LSP in ReopList : Yes
LSP ReopTime    : 4 minutes (Offset 15 seconds)
Next Reop Due in: 4 minutes 11 seconds
Exclude Path Constraint :
    10.0.1.7      exclude link
Computed ERO    :
    10.0.1.3
    10.0.1.5

PE1#
PE1#

```

Implementation Examples

Here is an example scenario and a solution for implementing RSVP LSP re-optimization monitoring.

- **Scenario 1:** RSVP tunnels are designed to take the shortest constraint-based path. When the network experiences a link or node failure, an established Label Switched Path (LSP) may transition to the next best available path computed at that time. However, once the failed network resources are restored or the network is expanded with new, more efficient links, the LSP remains on the suboptimal path. To ensure the LSP migrates back to the most efficient shorter path, the system requires a mechanism to trigger path recalculation.
- **Usecase 1:** By configuring the `reoptimize` command under the RSVP trunk or auto-bypass attributes, LSPs are added to a re-optimization wheel timer managed by the CSPF server. After the timer expires (default is 5 minutes), the system evaluates the LSP for re-optimization eligibility. If a new, shorter path is detected due to network restoration or expansion, the system seamlessly re-calculates and re-optimizes the LSP to the improved path. This ensures deterministic forwarding and maximizes resource utilization without manual intervention.

Revised Commands

Below is the revised command.

- `show cspf lsp` - for more details, refer to the ISIS Show Commands chapter in *Layer 3 guide*.
- [rsvp-session force-reoptimize \(page 952\)](#)

Troubleshooting

- **Timer Not Displayed:** Ensure the `reoptimize` command is configured under the specific `rsvp-trunk` or `rsvp-bypass` or `auto-bypass` attributes.
- **LSP Name Missing in CSPF:** Ensure the RSVP session is active and signaled; the name is synchronized from RSVP to the CSPF server.
- **Re-optimization Not Triggered:** Verify the "LSP in ReopList" field is set to "Yes" in the `show cspf lsp` output. If it is "No", the LSP is not currently eligible for automated re-optimization.

Glossary

| Key Terms/Acronym | Description |
|-------------------|---|
| CSPF | Constraint-based Shortest Path First; an extension to IGP protocols used to calculate paths for MPLS LSPs based on constraints. |
| LSP | Label Switched Path; a path through an MPLS network, set up by a signaling protocol such as RSVP. |
| RSVP | Resource Reservation Protocol; an MPLS signaling protocol that establishes label-switched paths (LSPs) with traffic engineering attributes. |
| Re-optimization | The process of recalculating an active LSP to find a more optimal path through the network. |
| Wheel Timer | A periodic timer mechanism used by the CSPF server to manage multiple LSPs and schedule their re-optimization checks. |

RSVP-TE COMMAND REFERENCE

| | |
|--------------------------------------|-----|
| RSVP-TE Commands | 869 |
| A.B.C.D | 872 |
| ack-wait-timeout | 873 |
| clear rsvp session | 874 |
| clear rsvp neighbor | 875 |
| clear rsvp trunk | 876 |
| cspf | 878 |
| debug rsvp all | 879 |
| debug rsvp cspf | 880 |
| debug rsvp events | 881 |
| debug rsvp fsm | 882 |
| debug rsvp hexdump | 883 |
| debug rsvp nsm | 884 |
| debug rsvp packet | 885 |
| description (rsvp-bypass) | 886 |
| description (rsvp-path) | 887 |
| description (rsvp-trunk) | 888 |
| disable-rsvp | 889 |
| enable-rsvp | 890 |
| entropy-label-capability | 891 |
| explicit-null | 892 |
| ext-tunnel-id A.B.C.D | 893 |
| from A.B.C.D | 894 |
| graceful-restart | 895 |
| graceful-restart recovery-time | 896 |
| graceful-restart restart-time | 897 |
| hello-interval | 898 |
| hello-receipt | 899 |
| hello-timeout | 900 |
| keep-multiplier | 901 |
| loop-detection | 902 |
| lsp-metric | 903 |
| lsp-reoptimization-timer | 904 |
| map-route A.B.C.D | 905 |
| neighbor A.B.C.D | 906 |
| neighbor X:X::X:X | 907 |
| no cspf | 908 |

| | |
|----------------------------------|-----|
| no loop-detection | 909 |
| no php | 910 |
| no primary | 911 |
| no primary affinity | 912 |
| no primary cspf | 913 |
| no primary record | 914 |
| no record | 915 |
| no refresh-path-parsing | 916 |
| no refresh-resv-parsing | 917 |
| path-option dynamic pce | 918 |
| php | 919 |
| primary ADMIN-GROUP-NAME | 920 |
| primary affinity | 922 |
| primary bandwidth | 923 |
| primary cspf | 924 |
| primary cspf-retry-limit | 925 |
| primary cspf-retry-timer | 926 |
| primary filter | 927 |
| primary hold-priority | 928 |
| primary hop-limit | 929 |
| primary label-record | 930 |
| primary path | 931 |
| primary record | 932 |
| primary retry-limit | 933 |
| primary retry-timer | 934 |
| primary reuse-route-record | 935 |
| primary setup-priority | 936 |
| primary traffic | 937 |
| refresh-time | 938 |
| refresh-path-parsing | 939 |
| refresh-resv-parsing | 940 |
| reoptimize | 941 |
| restart rsvp graceful | 942 |
| revert-timer | 943 |
| router rsvp | 944 |
| rsvp hello-interval | 945 |
| rsvp hello-receipt | 946 |
| rsvp hello-timeout | 947 |
| rsvp keep-multiplier | 948 |
| rsvp refresh-time | 949 |

| | |
|---|-----|
| rsvp-path | 950 |
| rsvp-trunk | 951 |
| rsvp-session force-reoptimize | 952 |
| rsvp-trunk force-switchover-secondary | 953 |
| rsvp-trunk-restart | 954 |
| secondary | 955 |
| secondary ADMIN-GROUP-NAME | 956 |
| secondary bandwidth | 958 |
| secondary cspf | 959 |
| secondary cspf-retry-limit | 960 |
| secondary cspf-retry-timer | 961 |
| secondary filter | 962 |
| secondary hold-priority | 963 |
| secondary hop-limit | 964 |
| secondary label-record | 965 |
| no secondary affinity | 966 |
| no secondary cspf | 967 |
| no secondary record | 968 |
| secondary path | 969 |
| secondary-priority path | 970 |
| secondary-priority hold-priority | 971 |
| secondary-priority setup-priority | 972 |
| secondary-priority label-record | 973 |
| secondary-priority hop-limit | 974 |
| secondary-priority bandwidth | 975 |
| secondary record | 976 |
| secondary retry-limit | 977 |
| secondary retry-timer | 978 |
| secondary reuse-route-record | 979 |
| secondary setup-priority | 980 |
| secondary traffic | 981 |
| session state timeout | 982 |
| shutdown | 983 |
| snmp restart rsvp | 984 |
| to A.B.C.D | 985 |
| update-type | 986 |
| Fast Reroute Commands | 987 |
| detour-allow-primary-upstream-path | 988 |
| default-frr-protection | 989 |
| detour-identification | 990 |

| | |
|---|------|
| primary fast-reroute admin-group-name | 991 |
| primary fast-reroute bandwidth | 992 |
| primary fast-reroute hold-priority | 993 |
| primary fast-reroute hop-limit | 994 |
| primary fast-reroute node-protection | 995 |
| primary fast-reroute protection | 996 |
| primary fast-reroute setup-priority | 997 |
| Refresh Reduction Commands | 998 |
| ack-send-timer | 999 |
| ack-wait-timeout | 1000 |
| message-ack | 1001 |
| refresh-reduction | 1002 |
| rsvp ack-wait-timeout | 1003 |
| rsvp message-ack | 1004 |
| rsvp refresh-reduction | 1005 |
| Facility Backup Commands | 1006 |
| backup-bw-type | 1007 |
| bandwidth | 1008 |
| bypass-lsp-addr-query-interval | 1009 |
| cspf-retry-limit | 1010 |
| cspf-retry-timer | 1011 |
| filter | 1012 |
| hold-priority | 1013 |
| hop-limit | 1014 |
| label-record | 1015 |
| no record | 1016 |
| path | 1017 |
| preemption-type | 1018 |
| record | 1019 |
| retry-limit | 1020 |
| retry-timer | 1021 |
| reuse-route-record | 1022 |
| rsvp-bypass | 1023 |
| setup-priority | 1024 |
| shutdown | 1025 |
| to A.B.C.D | 1026 |
| traffic | 1027 |
| Differentiated Services Commands | 1028 |
| map-route A.B.C.D | 1029 |
| map-route X:X::X:X | 1030 |

| | |
|--|------|
| override-diffserv | 1031 |
| primary class-to-exp-bit | 1032 |
| primary elsp-signaled | 1033 |
| primary llsp | 1034 |
| secondary elsp-signaled | 1035 |
| secondary llsp | 1036 |
| secondary map class | 1037 |
| show rsvp diffserv-info | 1038 |
| RSVP Multipath Commands | 1039 |
| description | 1040 |
| multipath-group | 1041 |
| rsvp-multipath | 1042 |
| to A.B.C.D | 1043 |
| show mpls counters rsvp multipath-name | 1044 |
| show mpls rsvp-multipath | 1045 |
| show rsvp multipath | 1046 |
| show running-config rsvp-multipath | 1047 |
| RSVP-TE Show Commands | 1048 |
| show rsvp label-pool | 1049 |
| show debugging rsvp | 1050 |
| show rsvp | 1051 |
| show rsvp admin-groups | 1053 |
| show rsvp bypass | 1054 |
| show rsvp bypass detail | 1055 |
| show rsvp bypass lsp-address-list | 1057 |
| show rsvp bypass protected-lsp-list | 1058 |
| show rsvp control-adjacency | 1060 |
| show rsvp data-link | 1062 |
| show rsvp graceful-restart | 1063 |
| show rsvp interface | 1064 |
| show rsvp l2-info | 1066 |
| show rsvp local-addresses | 1067 |
| show rsvp lsp-id-table | 1068 |
| show rsvp neighbor | 1070 |
| show rsvp nexthop-cache | 1072 |
| show rsvp path | 1073 |
| show rsvp protected-lsp-reop-list | 1074 |
| show rsvp session | 1075 |
| show rsvp session count | 1080 |
| show rsvp session egress | 1081 |

| | |
|--|------|
| show rsvp session ingress | 1085 |
| show rsvp session LSP-NAME | 1089 |
| show rsvp session transit | 1092 |
| show rsvp statistics | 1095 |
| show rsvp summary-refresh | 1096 |
| show rsvp trunk | 1097 |
| show rsvp trunk multi-sec-detail | 1099 |
| show rsvp trunk-id-table | 1100 |
| show rsvp version | 1102 |

RSVP-TE Commands

This chapter describes the RSVP-TE commands.

| | |
|--------------------------------------|-----|
| A.B.C.D | 872 |
| ack-wait-timeout | 873 |
| clear rsvp session | 874 |
| clear rsvp neighbor | 875 |
| clear rsvp trunk | 876 |
| cspf | 878 |
| debug rsvp all | 879 |
| debug rsvp cspf | 880 |
| debug rsvp events | 881 |
| debug rsvp fsm | 882 |
| debug rsvp hexdump | 883 |
| debug rsvp nsm | 884 |
| debug rsvp packet | 885 |
| description (rsvp-bypass) | 886 |
| description (rsvp-path) | 887 |
| description (rsvp-trunk) | 888 |
| disable-rsvp | 889 |
| enable-rsvp | 890 |
| entropy-label-capability | 891 |
| explicit-null | 892 |
| ext-tunnel-id A.B.C.D | 893 |
| from A.B.C.D | 894 |
| graceful-restart | 895 |
| graceful-restart recovery-time | 896 |
| graceful-restart restart-time | 897 |
| hello-interval | 898 |
| hello-receipt | 899 |
| hello-timeout | 900 |
| keep-multiplier | 901 |
| loop-detection | 902 |
| lsp-metric | 903 |
| lsp-reoptimization-timer | 904 |
| map-route A.B.C.D | 905 |
| neighbor A.B.C.D | 906 |
| neighbor X:X::X:X | 907 |
| no cspf | 908 |

| | |
|----------------------------------|-----|
| no loop-detection | 909 |
| no php | 910 |
| no primary | 911 |
| no primary affinity | 912 |
| no primary cspf | 913 |
| no primary record | 914 |
| no record | 915 |
| no refresh-path-parsing | 916 |
| no refresh-resv-parsing | 917 |
| path-option dynamic pce | 918 |
| php | 919 |
| primary ADMIN-GROUP-NAME | 920 |
| primary affinity | 922 |
| primary bandwidth | 923 |
| primary cspf | 924 |
| primary cspf-retry-limit | 925 |
| primary cspf-retry-timer | 926 |
| primary filter | 927 |
| primary hold-priority | 928 |
| primary hop-limit | 929 |
| primary label-record | 930 |
| primary path | 931 |
| primary record | 932 |
| primary retry-limit | 933 |
| primary retry-timer | 934 |
| primary reuse-route-record | 935 |
| primary setup-priority | 936 |
| primary traffic | 937 |
| refresh-time | 938 |
| refresh-path-parsing | 939 |
| refresh-resv-parsing | 940 |
| reoptimize | 941 |
| restart rsvp graceful | 942 |
| revert-timer | 943 |
| router rsvp | 944 |
| rsvp hello-interval | 945 |
| rsvp hello-receipt | 946 |
| rsvp hello-timeout | 947 |
| rsvp keep-multiplier | 948 |
| rsvp refresh-time | 949 |

| | |
|---|-----|
| rsvp-path | 950 |
| rsvp-trunk | 951 |
| rsvp-session force-reoptimize | 952 |
| rsvp-trunk force-switchover-secondary | 953 |
| rsvp-trunk-restart | 954 |
| secondary | 955 |
| secondary ADMIN-GROUP-NAME | 956 |
| secondary bandwidth | 958 |
| secondary cspf | 959 |
| secondary cspf-retry-limit | 960 |
| secondary cspf-retry-timer | 961 |
| secondary filter | 962 |
| secondary hold-priority | 963 |
| secondary hop-limit | 964 |
| secondary label-record | 965 |
| no secondary affinity | 966 |
| no secondary cspf | 967 |
| no secondary record | 968 |
| secondary path | 969 |
| secondary-priority path | 970 |
| secondary-priority hold-priority | 971 |
| secondary-priority setup-priority | 972 |
| secondary-priority label-record | 973 |
| secondary-priority hop-limit | 974 |
| secondary-priority bandwidth | 975 |
| secondary record | 976 |
| secondary retry-limit | 977 |
| secondary retry-timer | 978 |
| secondary reuse-route-record | 979 |
| secondary setup-priority | 980 |
| secondary traffic | 981 |
| session state timeout | 982 |
| shutdown | 983 |
| snmp restart rsvp | 984 |
| to A.B.C.D | 985 |
| update-type | 986 |

A.B.C.D

Use this command to configure an explicit IPv4 route sub-object as either loose or strict. A list of sub-objects specifies an explicit route to the egress router for an LSP.

- For the strict type of route addresses, the route taken from the previous router to the current router must be a directly connected path, and a message exchanged between the two routers should not pass any intermediate routers. This ensures that routing is enforced on the basis of each link. To configure the path option as a “strict” path, configure the interface IP address (IP address of the next-hop interface). Do not configure the “strict” path with loopback IP when the CSPF server is OSPF.
- For the loose type of route addresses, the route taken from the previous router to the current router does not need to be a direct path, and a message exchanged between the two routers can pass through other routers. To configure the path option as a “loose” path, configure either a loopback interface IP address or a next-hop interface IP address. This will expand the explicit path accordingly since the next hop need not be a connected one.

Use the `no` parameter with this command to disable the configuration.

Command Syntax

```
A.B.C.D
A.B.C.D (loose|strict)
no A.B.C.D
no A.B.C.D (loose|strict)
```

Parameters

loose

Make this node loose

strict

Make this node strict

Default

By default, A.B.C.D is disabled

Command Mode

Path mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-path mypath
(config-path)#10.10.0.5 strict
```

ack-wait-timeout

Use this command to set the acknowledgement wait timeout for the RSVP daemon. This command can be invoked from config-router mode.

Use the `no` parameter with this command to revert to the default settings.

Command Syntax

```
ack-wait-timeout <1-65535>
no ack-wait-timeout
```

Parameters

<1-65535>

Configure acknowledgment wait timeout.

Default

By default, `no ack-wait-timeout` is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 4.2.

Examples

```
#configure terminal
(config)router rsvp
(config-router)# ack-wait-timeout 30

#configure terminal
(config)router rsvp
(config-router)# no ack-wait-timeout
```

clear rsvp session

Use this command to reset either all or specified sessions originating from a specific ingress and terminating on the specific egress.



Note: If the affected session originates from the router where the command is issued, it is stopped and started. If the affected session does not originate from the router where the command is issued, it is stopped and deleted.

Command Syntax

```
clear rsvp session TUNNEL-ID LSP-ID INGRESS EGRESS
```

Parameters

TUNNELID

Clear tunnel ID sessions

LSP-ID

Clear LSP ID sessions

INGRESS

Clear ingress sessions

EGRESS

Clear egress sessions

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear rsvp session 1 1 1.2.3.4 192.168.1.1
```

clear rsvp neighbor

Use this command clears the RSVP (Resource Reservation Protocol) session with a specific neighbor. It is typically used to reset the RSVP signaling state with a peer router.

Command Syntax

```
clear rsvp neighbor A.B.C.D
```

Parameters

A.B.C.D

Use this parameter to clear the IP address of the IPv4 RSVP neighbor.

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 6.6.1.

Examples

```
#clear rsvp neighbor 10.1.4.14
```

clear rsvp trunk

Use this command to clear an RSVP trunk or to clear all RSVP trunks.

Clearing a trunk also kills any session associated with the trunk. This command is useful when a trunk is missing required data such as routing information. When data is missing, the trunk is in an incomplete state, and clearing it correctly re-initializes the session.



Note: If this command is given in the session on the ingress router, the session stops and restarts. If this command is given in the session on the egress router, the session is not cleared.

Command Syntax



Note: Use the following commands to clear standard RSVP Trunks:

```
clear rsvp trunk *
clear rsvp trunk ingress (TRUNKNAME|*)
clear rsvp trunk non-ingress (TRUNKNAME|*)
clear rsvp trunk (TRUNKNAME|*)
clear rsvp trunk (TRUNKNAME|*) primary
clear rsvp trunk (TRUNKNAME|*) secondary
```

Parameters

*

Clear all RSVP trunks configured

TRUNKNAME

Name of a specific trunk to be cleared

ingress

Clear an RSVP ingress trunk

non-ingress

Clear an RSVP non-Ingress trunk

primary

Clear all primary sessions configured for this trunk

secondary

Clear all secondary sessions configured for this trunk

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear rsvp trunk mytrunk
#clear rsvp trunk *
#clear rsvp trunk ingress mytrunk
#clear rsvp trunk ingress *
#clear rsvp trunk non-ingress mytrunk
#clear rsvp trunk non-ingress *
#clear rsvp trunk mytrunk primary
#clear rsvp trunk * primary
#clear rsvp trunk mytrunk secondary
#clear rsvp trunk * secondary
```

cspf

Use this command to enable the use of Constrained Shortest Path First (CSPF) server for all RSVP sessions. If CSPF is turned off globally, it cannot be enabled for any LSP.

The CSPF server computes paths for LSPs that are subject to various constraints such as bandwidth, hop count, administrative groups, priority, and explicit routes. When computing paths for LSPs, CSPF considers not only the topology of the network and the attributes defined for the LSP but also the links. It attempts to minimize congestion by intelligently balancing the network load.

Use the [no cspf \(page 908\)](#) command to disable this configuration.



Notes:

- CSPF server information is not signaled across session and hence sessions in transit and egress nodes will not be aware of the CSPF server. So, in multi CSPF scenarios, neighbor down event from a CSPF server restart all sessions irrespective of which CSPF server sessions were using.
- The CSPF server is unaware of whether an interface is label-switching and RSVP enabled. So, the LSP calculation may include non-RSVP interfaces too, and may lead to tunnel sessions not coming up. To avoid links that are not label-switching and RSVP enabled, “te-metric” with the highest value shall be configured on those interfaces to de-prioritize those links in the TE database.

Command Syntax

```
cspf
```

Parameters

None

Default

By default, CSPF server is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows using the `no cspf` command in Router mode to disable CSPF for all RSVP sessions.

```
#configure terminal
(config)#router rsvp
(config-router)#cspf
```

debug rsvp all

Use this command to enable all debugging options for an RSVP daemon.

Use the `no` parameter with this command to stop logging all debugging information.

Command Syntax

```
debug rsvp (all|)  
no debug rsvp (all|)
```

Parameters

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug rsvp all
```

debug rsvp cspf

Use this command to enable the exchange of debugging messages between the RSVP module and the CSPF module.

Use the `no` parameter with this command to stop logging CSPF debugging information.

Command Syntax

```
debug rsvp cspf  
no debug rsvp cspf
```

Parameters

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug rsvp cspf
```

debug rsvp events

Use this command to enable debugging of events that were generated from an RSVP daemon.

Use the `no` parameter with this command to stop logging RSVP debugging information.

Command Syntax

```
debug rsvp events
no debug rsvp events
```

Parameters

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug rsvp events
```

debug rsvp fsm

Use these commands to enable debugging of events related to RSVP finite state machines (FSM). Commands are available to log debugging information for the egress FSM, the ingress FSM, the transit FSM, the transit upstream FSM, or the transit downstream FSM.

Use the `no` parameter with these commands to stop logging FSM debugging information.

Command Syntax

```
debug rsvp fsm
debug rsvp fsm egress
debug rsvp fsm ingress
debug rsvp fsm transit
debug rsvp fsm transit upstream
debug rsvp fsm transit downstream
no debug rsvp fsm
no debug rsvp fsm egress
no debug rsvp fsm ingress
no debug rsvp fsm transit
no debug rsvp fsm transit upstream
no debug rsvp fsm transit downstream
```

Parameters

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)#debug rsvp fsm transit upstream
```

debug rsvp hexdump

Use this command to enable the hexdump debugging option for an RSVP daemon.

Use the `no` parameter with this command to stop logging hexdump debugging information.

Command Syntax

```
debug rsvp hexdump
no debug rsvp hexdump
```

Parameters

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug rsvp hexdump
```

debug rsvp nsm

Use this command to enable the NSM debugging option for an RSVP daemon.

Use the `no` parameter with this command to stop logging NSM debugging information.

Command Syntax

```
debug rsvp nsm
no debug rsvp nsm
```

Parameters

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug rsvp nsm
```

debug rsvp packet

Use this command to enable packet debugging options for an RSVP daemon. Using the `in` option command enables debugging for incoming packets. Using the `out` option command enables debugging for outgoing packets. Use the `no` parameter with these commands to stop logging debugging information.

Command Syntax

```
debug rsvp packet
debug rsvp packet in
debug rsvp packet out
no debug rsvp packet
no debug rsvp packet in
no debug rsvp packet out
```

Parameters

None

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#debug rsvp packet in
#debug rsvp packet out
```

description (rsvp-bypass)

Use this command to add a description to the rsvp-bypass or update an existing description.

Use the `no` parameter with this command to disable the configuration.

Command Syntax

```
description LINE
no description
```

Parameters

LINE

Line describing the RSVP tunnel

Default

By default, rsvp bypass description is empty.

Command Mode

Rsvp-bypass mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#description this_is_the_description
(config-bypass)#no description
```

description (rsvp-path)

Use this command to add a description to the rsvp-path or update an existing description.

Use the `no` parameter with this command to disable the configuration.

Command Syntax

```
description LINE
no description
```

Parameters

LINE

Line describing the RSVP path

Default

By default, rsvp path description is empty.

Command Mode

Rsvp-path mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
#configure terminal
(config)#rsvp-path mypath
(config-path)#description this_is_the_description
(config-path)#no description
```

description (rsvp-trunk)

Use this command to add a description to the rsvp-trunk or update an existing description.

Use the `no` parameter with this command to disable the configuration.

Command Syntax

```
description LINE
no description
```

Parameters

LINE

Line describing the RSVP tunnel

Default

By default, rsvp trunk description is empty.

Command Mode

Rsvp-trunk mode.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#description this_is_the_description
(config-trunk)#no description
```

disable-rsvp

Use this command to disable RSVP message exchange on an interface.

RSVP can be enabled using the [enable-rsvp \(page 890\)](#) command.

Command Syntax

```
disable-rsvp
```

Parameters

None

Default

By default, RSVP message exchange is disabled on an interface.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#disable-rsvp
```

enable-rsvp

Use this command to enable RSVP message exchange on an interface.



Note: To use this command, the corresponding interface needs to be enabled for label-switching using the [label-switching \(page 1396\)](#) command.

See [disable-rsvp \(page 889\)](#) to undo the effects of this command.

Command Syntax

```
enable-rsvp
```

Parameters

None

Default

By default, RSVP message exchange is disabled.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth1
(config-if)#enable-rsvp
```

entropy-label-capability

Use this command to share the load across multiple members of a LAG port in the core of an MPLS network by using entropy labels.

Use the `no` form of the command to disable the use of entropy labels



Note: Load balancing is enabled by default for all the parameters. If you enable load balancing manually, then all the parameters enabled by default are reset and you need to enable the parameters based on which traffic should be load balanced.

Command Syntax

```
entropy-label-capability
no entropy-label-capability
```

Parameters

None

Default

By default, entropy labels are not used.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 3.0.

Examples

```
#configure terminal
(config)router rsvp
(config-router)#entropy-label-capability
```

explicit-null

Use this command to send explicit-null labels to upstream router, instead of implicit-null labels.

If php is enabled then implicit-null label is advertised, then the penultimate hop removes the label and sends the packet as a plain IP packet to the egress router. The explicit-null command advertises label 0 and retains the label so the egress router can pop it. For details about usage of explicit-null, please refer to RFC 3032.

Use the `no` parameter with this command to stop sending explicit-null labels for directly-connected FECs to upstream router and resume sending non reserved labels.

Command Syntax

```
explicit-null
no explicit-null
```

Parameters

None

Default

By default, no php is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#explicit-null
```


ext-tunnel-id A.B.C.D

Use this command to configure an extended-tunnel identifier as an IPv4 address. These identifiers are used in RSVP messages. If no extended-tunnel ID is specified, the LSR-ID for the router is used as the extended-tunnel ID for all LSPs. The extended-tunnel ID is a simple way of identifying all LSPs belonging to the same trunk.

Use the no parameter with this command to remove a configured extended-tunnel ID.

Command Syntax

```
ext-tunnel-id A.B.C.D
no ext-tunnel-id
```

Parameters

A.B.C.D

Extended tunnel identifier for this trunk in IPv4 address format

Default

By default, the LSR-ID of the router is used as the extended-tunnel ID for all sessions.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk t1
(config-trunk)#ext-tunnel-id 10.10.10.30

(config)#rsvp-trunk t1
(config-trunk)#no ext-tunnel-id
```

from A.B.C.D

Use this command to specify a “from” IPv4 address for the RSVP daemon. This command can be invoked from either the [router rsvp \(page 944\)](#) mode or from the [rsvp-trunk \(page 951\)](#) mode. In the RSVP router mode, this command defines the source address as an IPv4 packet sent out by the RSVP daemon. In the RSVP trunk mode, this command indicates a sender’s address in the sender template object that is used in path messages.

Use the `no` parameter with this command to revert to the default settings.

Command Syntax

```
from A.B.C.D
no from
```

Parameters

A.B.C.D

When in trunk mode, this is the IPv4 address of a tunnel ingress node

A.B.C.D

When in router mode, this is the loopback IPv4 address

Default

By default, from A.B.C.D is enabled

Command Mode

Router or Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#from 10.10.0.5

#configure terminal
(config)#router rsvp
(config-router)#from 10.10.0.5
```

graceful-restart

Use this command to enable RSVP-TE graceful restart capability on a router. This is a global parameter. RSVP-TE determines whether or not to send the graceful restart capability object in its hello message. However, this capability also depends on support for graceful restart on the neighboring router.



Notes:

The following conditions must be met to activate RSVP-TE graceful restart:

- This command is used on the local router.
- The neighbor router must be explicitly configured with either the [entropy-label-capability \(page 891\)](#) or the [neighbor X:X::X:X \(page 907\)](#) command.
- The neighbor router supports graceful restart, and it is activated.
- Graceful restart is supported only for primary and secondary sessions; it is not supported for FRR or multiple-secondary entities.

Command Syntax

```
graceful-restart
no graceful-restart
```

Parameters

None

Default

Disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 5.0.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#graceful-restart
(config-router)#no graceful-restart
```

graceful-restart recovery-time

Use this command to set a recovery time for an RSVP-TE graceful restart configuration.

Use the `no` parameter with this command to reset the recovery time.

Command Syntax

```
graceful-restart recovery-time <60000-3600000>  
no graceful-restart recovery-time
```

Parameters

<60000-3600000>

Recovery time value in milliseconds

Default

Default value is 360000 ms.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 5.0.

Examples

```
#configure terminal  
(config)#router rsvp  
(config-router)#graceful-restart recovery-time 600000
```

graceful-restart restart-time

Use this command to set a restart time for an RSVP-TE graceful restart configuration.

Use the `no` parameter with this command to reset the restart time.

Command Syntax

```
graceful-restart restart-time <10000-600000>  
no graceful-restart restart-time
```

Parameters

<10000-600000>

Restart time value in milliseconds

Default

Default value is 180000 ms.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 5.0.

Examples

```
#configure terminal  
(config)#router rsvp  
(config-router)#graceful-restart restart-time 100000
```

hello-interval

Use this command to set an interval between Hello packets.

Used as a global command, this value is over-ridden by the hello-interval set on the interface (see [rsvp hello-interval \(page 945\)](#)). For optimum performance, set this value no more than one-third of the hello-timeout value.

Use the `no` parameter with this command to return to the default hello interval value.

Command Syntax

```
hello-interval <1-65535>
no hello-interval
```

Parameter

<1-65535>

The time in seconds after which hello packets are sent

Default

By default, hello interval is 2 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#hello-interval 5

(config)#router rsvp
(config-router)#no hello-interval
```

hello-receipt

Use this command to enable the receipt of Hello messages from peers.

Use the `no` parameter with this command to disable the exchange of Hello messages.

Command Syntax

```
hello-receipt
no hello-receipt
```

Parameters

None

Default

By default, hello receipt is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#hello-receipt
```

hello-timeout

Use this command to set the RSVP hello timeout. If an LSR has not received a hello message from a peer within the number of seconds set with this command, all sessions shared with this peer are reset. The hello-timeout determines how long an RSVP node waits for a hello message before declaring a neighbor to be down.

Use the `no` parameter with this command to set the default hello timeout value.

Command Syntax

```
hello-timeout <1-65535>
no hello-timeout
```

Parameter

<1-65535>

Time in seconds to receive a hello message.

Default

By default, hello-timeout value is 7 seconds.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#hello-timeout 12

(config)#router rsvp
(config-router)#no hello-timeout
```


keep-multiplier

Use this command to configure the constant to be used to calculate a valid reservation lifetime for a Labeled Switched Path (LSP).

The refresh time and keep multiplier are two interrelated timing parameters used to calculate the valid reservation lifetime for an LSP. Use the following formula to calculate the reservation lifetime for an LSP:

$$L \geq (K + 0.5) * 1.5 * R$$

K = keep-multiplier

R = refresh timer

The router sends refresh messages periodically so that the neighbors do not timeout.

Use the `no` parameter with this command to return to the default keep-multiplier setting.

Command Syntax

```
keep-multiplier <1-255>
no keep-multiplier <1-255>
```

Parameters

<1-255>

The keep-multiplier value

Default

By default, keep-multiplier value is 3

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#keep-multiplier 2
```

loop-detection

Use this command to turn on loop detection for Path and Reservation messages exchanged between LSRs.

Use the [no loop-detection \(page 909\)](#) command to return to default settings.

Command Syntax

```
loop-detection
```

Parameters

None

Default

By default, loop detection is enabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#loop-detection
```

lsp-metric

Use this command to set LSP absolute metric or relative metric for IGP Shortcut use

Use the `no` parameter along with this `command` to unset the LSP metric for IGP shortcut.

Command Syntax

```
lsp-metric absolute <1-65535>
lsp-metric relative (<-65535-0>|<1-65535>)
no lsp-metric absolute (<1-65535>|)
no lsp-metric relative (<-65535-0>|<1-65535>|)
```

Parameters

absolute

Absolute metric

relative

Relative metric

<1-65535>

Metric value

<-65535-0>

The keep-multiplier value

Command Mode

RSVP Trunk mode

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#exit
(config)#rsvp-trunk T1
(config-trunk)#lsp-metric absolute 10
(config-trunk)#lsp-metric relative 10
```

lsp-reoptimization-timer

Use this command to set the re-optimization interval timer.

Use the no parameter with this command to set the default re-optimization interval (5 minutes).

Command Syntax

```
lsp-reoptimization-timer <1-240>
```

Parameter

<1-240>

The interval in minutes after which LSP re-optimization will take place.

Default

By default, the re-optimization timer interval is 5 minutes.

Command Mode

RSVP router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config)#router rsvp
(config-router)#lsp-reoptimization-timer 10
(config)#router rsvp
(config-router)#no lsp-reoptimization-timer
```

map-route A.B.C.D

Use this command to map a route using an IPv4 to an RSVP trunk. If the primary LSP for a trunk goes down, all mapped routes are sent automatically to a secondary LSP configured as backup for a primary LSP.

Use the no parameter with this command to unmap routes from specified trunks.

Command Syntax

```
map-route A.B.C.D/M
map-route A.B.C.D/M CLASS
map-route A.B.C.D A.B.C.D
map-route A.B.C.D A.B.C.D CLASS
no map-route A.B.C.D/M
no map-route A.B.C.D/M CLASS
no map-route A.B.C.D A.B.C.D
no map-route A.B.C.D A.B.C.D CLASS
```

Parameters

A.B.C.D/M

Prefix to map, plus mask

A.B.C.D

Prefix to be mapped

A.B.C.D

Prefix mask

CLASS

Incoming DiffServ Class (for example, be, ef, etc.) to map to the RSVP trunk

Default

By default, map route A.B.C.D/M is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#map-route 2.2.2.2/16
```

neighbor A.B.C.D

Use this command to designate a neighbor IPv4 address to use when exchanging hello messages. Any neighbor hello message that is not explicitly identified is rejected.

Use the `no` parameter with this command to remove an IP neighbor from the system.

Command Syntax

```
neighbor A.B.C.D  
no neighbor A.B.C.D
```

Parameters

None

Default

By default, neighbor A.B.C.D is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#router rsvp  
(config-router)#neighbor 10.10.0.5
```

neighbor X:X::X:X

Use this command to designate a neighbor IPv6 address to use when exchanging hello messages. Any neighbor hello message that is not explicitly identified is rejected.

Use the `no` parameter with this command to remove an IP neighbor from the system.

Command Syntax

```
neighbor X:X::X:X  
no neighbor X:X::X:X
```

Parameters

None

Default

By default, neighbor X:X::X:X is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#router rsvp  
(config-router)#neighbor 3ffe::3:34
```

no cspf

Use this command to disable the use of the Constrained Shortest Path First (CSPF) server for all RSVP sessions. Disable CSPF when no nodes support the required traffic engineering extensions.

When this command is executed in Router mode, CSPF is disabled for all configured RSVP sessions, and all RSVP sessions configured from this point forward. If the default CSPF per RSVP session is enabled, it will be disabled. The CSPF status for RSVP sessions can be verified using the [show rsvp session \(page 1075\)](#) command with the detail option.

Use the [cspf \(page 878\)](#) command to revert to the default settings.



Note: When CSPF is disabled, path is not calculated taking constraints into consideration. Path message is sent to the next hop based on IGP best route. In this case, ERO is not included in path message and all constraints are included.

Command Syntax

```
no cspf
```

Parameters

None

Default

By default, no cspf is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows using the `no cspf` command in Router mode to disable CSPF for all RSVP sessions.

```
#configure terminal
(config)#router rsvp
(config-router)#no cspf
```


no loop-detection

Use this command to turn off loop detection for Path and Reservation messages exchanged between LSRs. When a Path or Resv message is received, the primary IP address of the incoming interface is compared with the received route record list.

Use the [loop-detection \(page 902\)](#) command to revert to default settings.

Command Syntax

```
no loop-detection
```

Parameters

None

Default

By default, no loop detection is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#no loop-detection
```

no php

Use this command to disable Penultimate-Hop-Popping (PHP) for the router. An egress router sends neither implicit null label nor explicit null for LSPs. When `no php` command is used, the egress router sends non-reserved labels (those labels in the label pool range allotted to RSVP) to the upstream router and retains the labels till the egress router.



Note: Use the [show rsvp \(page 1051\)](#) command to display the status of Penultimate-Hop-Popping.

Command Syntax

```
no php
```

Parameters

None

Default

By default, no php is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#no php
```

no primary

Use this command to reset RSVP primary session to default configuration.

Command Syntax

```
no primary
```

Parameters

None

Default

None

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
Example to delete primary tunnel.  
  
(config)#rsvp-trunk mytrunk  
(config-trunk)#primary path mypath  
(config-trunk)#no primary
```

no primary affinity

Use this command to disable the use of sending out session attribute objects with resource affinity data.

Use the [primary affinity \(page 922\)](#) command to return to the default settings.

Command Syntax

```
no primary affinity
```

Parameters

None

Default

By default, primary no affinity is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#no primary affinity
```

no primary cspf

Use this command to disable the use of Constrained Shortest Path First (CSPF) server for an explicit route to the egress, or all RSVP sessions. When CSPF is turned off globally it cannot be enabled for any LSP. If used per LSP, it can be used to turn off CSPF computation for a specific LSP.

Disable CSPF when all nodes do not support the required traffic engineering extensions, and configure LSPs manually to use an explicit path. The LSP is then established only along the path specified by the operator.

Use the [primary cspf \(page 924\)](#) command to enable this setting.

Command Syntax

```
no primary cspf
```

Parameters

None

Default

By default, no primary cspf is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows using the `no cspf` command in Trunk mode to disable CSPF for the primary LSP.

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#no primary cspf
```

no primary record

Use this command to disable recording of the route taken by Path and Reservation Request (Resv) messages to confirm establishment of reservations and identify errors. Routes are recorded by means of the Route Record Object (RRO) in RSVP messages.

Use the [primary record \(page 932\)](#) command to return to the default settings.

Command Syntax

```
no primary record
```

Parameters

None

Default

By default, routes are recorded

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#no primary record
```

no record

Use this command to disable recording of the route taken by Path and Reservation Request (Resv) messages that confirm establishment of reservations and are used to identify errors. The routes are recorded by means of the Route Record Object (RRO) in RSVP messages.

Command Syntax

```
no record
```

Parameters

None

Default

Routes are recorded by default.

Command Mode

RSVP Bypass mode

Examples

```
#configure terminal
(config)#rsvp-bypass bypassname
(config-bypass)#no record
```

no refresh-path-parsing

Use this command to disable parsing of Refresh PATH messages received from upstream nodes. Enable this command to minimize message processing by RSVP, if you are sure that a particular router does not need to parse Refresh-PATH messages to check for changes because LSPs passing through this router are not required to be updated, simultaneously.

Use the [refresh-path-parsing \(page 939\)](#) command to revert to the default settings.

Command Syntax

```
no refresh-path-parsing
```

Parameters

None

Default

By default, refresh-path-parsing is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
Router#configure terminal
Router(config)#router rsvp
Router(config-router)#no refresh-path-parsing
```

no refresh-resv-parsing

Use this command to disable parsing of Refresh RESV messages received from upstream nodes. Enable this command to minimize message processing by RSVP, if you are sure that a particular router does not need to parse Refresh RESV messages to check for changes because LSPs passing through this router are not required to be updated simultaneously.

Command Syntax

```
no refresh-resv-parsing
```

Parameters

None

Default

By default, refresh reservation parsing is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
Router#configure terminal
Router(config)#router rsvp
Router(config-router)#no refresh-resv-parsing
```

path-option dynamic pce

Use this command to indicate that a tunnel must query PCE for path computation.

Command Syntax

```
path-option dynamic pce
```

Parameters

None

Command Mode

RSVP trunk mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#configure terminal
(config)#rsvp-trunk mytrunk ipv4
(config-trunk)#path-option dynamic pce
(config-trunk)#pce entity 1
```

php

Use this command to enable Penultimate-Hop-Popping for the router. An egress router send an implicit-null label (3) to the upstream router.



Note: Use the [show rsvp \(page 1051\)](#) command to display the status of Penultimate-Hop-Popping.

Use the [no php \(page 910\)](#) command to revert to the default setting.



Note: On Qumran1 (QAX, QMX, and QUX) platforms, when a primary session has non-implicit-null out label and a backup session has implicit-null out label, then L3 services do not work when the session is at backup state. i.e., mpls ping, L3 ping, and map route traffic will fail as label popped packets cannot select the next header ethertype properly.

Command Syntax

```
php
```

Parameters

None

Default

By default, no php is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#php
```

primary ADMIN-GROUP-NAME

Use this command to configure primary administrative groups. Administrative groups are manually assigned attributes that describe the color of links, so that links with the same color are in one class. These groups are used to implement different policy-based LSP setups. Administrative group attributes can be included or excluded for an LSP or for a path's primary and secondary paths.



Note: A link can be added to a specific Administrative Group via the Network Services Module. Refer to the Network Services Module Command Reference for details.

Use the `no` parameter to remove a previously configured group from an administrative group list.

Command Syntax

```
primary (include-any|include-all|exclude-any) ADMIN-GROUP-NAME
primary (include-any|exclude-any) ADMIN-GROUP-NAME
primary (include-any|include-all|exclude-any) ADMIN-GROUP-NAME
primary (include-any|exclude-any) ADMIN-GROUP-NAME
```

Parameters

include-any

Include any attributes

include-all

Include all attributes

exclude-any

Exclude any attribute

ADMIN-GROUP-NAME

Administrative group name

Default

By default, primary admin group name is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary exclude-any myadmingroup

#configure terminal
(config)#rsvp-trunk mytrunk
```

```
(config-trunk)#primary include-all admingrp2

#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary include-any admingrp2
```

primary affinity

Use this command to enable sending of session attribute objects with resource affinity data.

Use the [no primary affinity \(page 912\)](#) command to disable sending of session attribute objects.

Command Syntax

```
primary affinity
```

Parameters

None

Default

By default, primary affinity is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary affinity
```

primary bandwidth

Use this command to reserve the primary bandwidth in bits per second for the current trunk.

Each LSP has an associated bandwidth attribute. The bandwidth value is included in the sender's RSVP Path message and specifies the bandwidth to be reserved for the LSP. It is specified in bits per second, with a higher value indicating a greater user traffic volume. A zero bandwidth reserves no resources, although exchanges labels.

Use the `no` parameter to remove configured bandwidth information.

Command Syntax

```
primary bandwidth BANDWIDTH  
no primary bandwidth BANDWIDTH
```

Parameter

BANDWIDTH

<1-999>k for 1 to 999 kilobits/s

<1-999>m for 1 to 999 megabits/s

<1-100>g for 1 to 100 gigabits/s

Default

The default bandwidth is 0 bits per second, which allows data to flow through but does not reserve bandwidth.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#primary bandwidth 100m
```

primary cspf

Use this command to enable the use of Constrained Shortest Path First (CSPF) server for an explicit route to the egress, or all RSVP sessions. When CSPF is turned off globally, it cannot be enabled for any LSP.

The CSPF server computes paths for LSPs that are subject to constraints such as bandwidth, hop count, administrative groups, priority, and explicit routes. When computing paths for LSPs, CSPF considers not only the topology of the network and the attributes defined for the LSP, but also the links. It attempts to minimize congestion by intelligently balancing the network load.

Use the [no primary affinity \(page 912\)](#) command to revert to the default settings.



Note: CSPF server information is not signaled across session and hence sessions in transit and egress nodes will not be aware of the CSPF server. So, in multi CSPF scenarios, neighbor down event from a CSPF server restart all sessions irrespective of which CSPF server sessions were using.

Command Syntax

```
primary cspf
```

Parameters

None

Default

By default, primary cspf is enabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary cspf
```

primary cspf-retry-limit

Use this command to specify the number of retries that CSPF should carry out for a request received from RSVP. Use the `no` parameter with this command to disable this configuration.

Command Syntax

```
primary cspf-retry-limit <1-65535>
no primary cspf-retry-limit
```

Parameter

<1-65535>

Set the number of times CSPF should retry for this LSP

Default

By default, retry-limit is 0.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#primary cspf-retry-limit 535

(config)#rsvp-trunk T1
(config-trunk)#no primary cspf-retry-limit
```

primary cspf-retry-timer

Use this command to specify the time between each retry that CSPF might carry out for a request received from RSVP.

Use the `no` parameter with this command to disable this configuration.

Command Syntax

```
primary cspf-retry-timer <1-600>  
no primary cspf-retry-timer
```

Parameter

<1-600>

Timeout between successive retries, in seconds

Default

By default, retry-timer is 0

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk T1  
(config-trunk)#primary cspf-retry-timer 45  
  
(config)#rsvp-trunk T1  
(config-trunk)#no primary cspf-retry-timer
```

primary filter

Use this command to set the filter to the fixed or shared style for an LSP.

- The shared filter style identifies a shared reservation environment. It creates a single reservation into which flows from all senders are mixed.
- The fixed filter style designates a distinct reservation. A distinct reservation request is created for data packets from a particular sender. The fixed filter style is also used style to prevent rerouting of an LSP and to prevent another LSP from using this bandwidth.

Use the `no` parameter to reset the configured filter to the default.

Command Syntax

```
primary filter fixed
no primary filter
```

Parameters

fixed

Use a fixed filter for this LSP

Default

By default, primary filter is shared-explicit.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary filter fixed
```

primary hold-priority

Use this command to configure the hold priority value for the selected trunk. In case of insufficient bandwidth, remove less important existing LSPs to free up a portion of the bandwidth. This can be done by preempting one or more of the signaled LSPs. Hold priority determines the degree to which an LSP holds onto its reservation for a session after the LSP has been configured successfully. When the hold priority is high, the existing LSP is less likely to give up its reservation.

Use the `no` parameter to reset the trunk to the default hold-priority value.

Command Syntax

```
primary hold-priority <0-7>
no primary hold-priority
```

Parameter

<0-7>

Set a hold priority for the LSP

Default

The default hold-priority value is 0, which is the highest. Once a session is configured with a hold priority of 0, no other session can preempt it.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary hold-priority 2
```

primary hop-limit

Use this command to specify a limit of hops for an RSVP trunk. Hop-limit data is sent to the CSPF server if CSPF is used.

Upon configuration of an arbitrary hop-limit, the hop-limit is compared with the number of hops configured in the primary path, if a primary path has been configured. If the number of hops in the primary path exceeds the hop-limit configured, no `Path` messages are sent, and any existing session is torn down. If no primary path is configured, the trunk is processed normally and `Path` messages are sent.

Use the `no` parameter to reset the trunk to the default hop-limit value.

Command Syntax

```
primary hop-limit <1-255>
no primary hop-limit <1-255>
no primary hop-limit
```

Parameters

<1-255>

Set the number of acceptable hops for the LSP

Default

By default, primary hop limit is 255.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary hop-limit 23
```

primary label-record

Use this command to record all labels exchanged between RSVP-enabled routers during the reservation setup process.

Use the `no` parameter with this command to turn off recording.

Command Syntax

```
primary label-record  
no primary label-record
```

Parameters

None

Default

By default, primary label record is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#primary label-record
```

primary path

Use this command to specify an RSVP path to be used. The `PATHNAME` in this command is the string (name) used to identify an RSVP path defined for the node (refer to `rsvp-path` command).

Use the `no` parameter with this command to remove a configured RSVP path.

Command Syntax

```
primary path  
PATHNAME  
no primary path
```

Parameters

PATHNAME

The name of the path to use

Default

By default, primary path is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#primary path mypath
```

primary record

Use this command to enable recording of the route taken by Path and Reservation Request (Resv) messages to confirm establishment of reservations and identify errors. Routes are recorded by means of the Route Record Object (RRO) in RSVP messages.

Use the [no primary record \(page 914\)](#) command to disable recording of routes.

Command Syntax

```
primary record
```

Parameters

None

Default

By default, routes are recorded

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary record
```

primary retry-limit

Use this command to specify a retry count this RSVP Trunk.

If a session is in a `nonexistent` state due to a path error message, the system tries to recreate the LSP for the number of times specified by the `retry-limit` command.

Although the same retry command controls both the trunk and the session, the retry-limit value affects only the session and not the trunk. If the trunk is in an `incomplete` state, the code keeps trying forever to bring it to a `complete` state regardless of the retry-limit value.

Use the `no` parameter with this command to revert to the default retry-limit value.

Command Syntax

```
primary retry-limit <1-65535>
no primary retry-limit
```

Parameter

<1-65535>

The set number of times the system should try setting up the LSP

Default

By default, the retry-limit value is 0, and the trunk and session try to create the LSP indefinitely.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary retry-limit 256
```

primary retry-timer

Use this command to specify a retry interval for an RSVP Trunk. When an ingress node tries to configure an LSP and the setup fails due to the receipt of a Path Error message, the system waits for the time configured with this command, before retrying the LSP setup process.

Use the `no` parameter with this command to revert to the default retry-time value.

Command Syntax

```
primary retry-timer <1-600>
no primary retry-timer
```

Parameter

<1-600>

Time in seconds after which the system should retry setting up the LSP

Default

By default, retry-timer value is 30 seconds.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary retry-timer 12
```

primary reuse-route-record

Use this command to use the updated Route Record List as an Explicit Route (with all strict nodes) when a path message is sent out at the next refresh.

The ERO list contains the hops to be taken to reach the egress from the current LSR. If CSPF is not available, to place an ERO with all strict routes, use this command to modify the ERO after receiving the Resv message. The future Path messages have the ERO with all strict nodes, identifying each and every node to be traversed.

Use the `no` parameter with this command to disable the use of the Route Record List as the explicit route.

Command Syntax

```
primary reuse-route-record  
no primary reuse-route-record
```

Parameters

None

Default

By default, primary reuse route record is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#primary reuse-route-record
```

primary setup-priority

Use this command to configure a setup priority value for a trunk. In case of insufficient bandwidth, users must remove less important LSPs to free up the bandwidth. This can be done by preempting one or more of the existing LSPs. The primary setup priority determines if a new LSP can preempt an existing LSP.

The setup priority of the new LSP must be higher than the hold priority of an existing LSP for the existing LSP to be preempted. Note that for a trunk, the setup priority should not be higher than the hold priority.

Use the `no` parameter with this command to revert to the default primary setup priority value.

Command Syntax

```
primary setup-priority <0-7>
no primary setup-priority
```

Parameters

<0-7>

Set the priority value

Default

By default, setup priority is 7, which is the lowest.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary setup-priority 2
```

primary traffic

Use this command to specify the traffic type for this RSVP Trunk.

Use the `no` parameter with this command to reset the configured traffic type.

Command Syntax

```
primary traffic (guaranteed|controlled-load)
no primary traffic
```

Parameters

controlled-load

Controlled loaded traffic

guaranteed

Guaranteed traffic

Default

By default, primary traffic type is controlled-load

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#primary traffic guaranteed
```

refresh-time

Use this command to configure RSVP refresh interval timer. The timer specifies the interval after which Path and/ or Reservation Request (Resv) messages will be sent out.

The refresh time and keep multiplier are two interrelated timing parameters used to calculate the valid Reservation Lifetime for an LSP. Refresh time regulates the interval between Refresh messages which include Path and Reservation Request (Resv) messages. Refresh messages are sent periodically so that reservation does not timeout in the neighboring nodes. Each sender and receiver host sends Path and Resv messages, downstream and upstream respectively, along the paths.

Use the `no` parameter with this command to return to the default refresh-time interval.

Command Syntax

```
refresh-time <1-65535>  
no refresh-time <1-65535>  
no refresh-time
```

Parameter

<1-65535>

The duration for which messages are sent, in seconds

Default

By default, refresh-time interval is 30 seconds

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#router rsvp  
(config-router)#refresh-time 20
```

refresh-path-parsing

Use this command to disable parsing of Refresh PATH messages received from upstream nodes. Use this command to minimize message processing by RSVP when you are sure that a particular router does not need to parse Refresh-PATH messages to check for changes, because LSPs passing through this router are not required to be updated simultaneously.

Use the [no refresh-path-parsing \(page 916\)](#) command to disable this setting.

Command Syntax

```
refresh-path-parsing
```

Parameters

None

Default

By default, refresh-path-parsing is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
Router#configure terminal
Router(config)#router rsvp
Router(config-router)#refresh-path-parsing
```

refresh-resv-parsing

Use this command to disable parsing of Refresh RESV messages received from upstream nodes. Use this command to minimize message processing by RSVP when you are sure that a particular router does not need to parse Refresh RESV messages to check for changes because LSPs passing through this router are not required to be updated simultaneously.

Use the [no refresh-resv-parsing \(page 917\)](#) command to disable this setting.

Command Syntax

```
refresh-resv-parsing
```

Parameters

None

Default

By default, refresh reservation parsing is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
Router#configure terminal
Router(config)#router rsvp
Router(config-router)#refresh-resv-parsing
```


reoptimize

Use this command to enable re-optimization of the RSVP primary LSP with the [rsvp-session force-reoptimize \(page 952\)](#) command.

Use the no parameter with this command to disable re-optimization of the RSVP primary LSP.

Command Syntax

```
reoptimize
```

Parameters

None

Command Mode

Trunk mode

Applicability

This command was introduced in OcNOS version 1.3.4.

Examples

```
(config)#rsvp-trunk t1
(config-trunk)#reoptimize
(config)#rsvp-trunk t1
(config-trunk)#no reoptimize
```

restart rsvp graceful

Use this command to restart RSVP gracefully.

To restart RSVP gracefully, you must give the [graceful-restart \(page 895\)](#) command to enable graceful restart capability on the device in RSVP router mode.

Command Syntax

```
restart rsvp graceful
```

Parameter

None

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 5.0.

Example

```
#restart rsvp graceful
#restart rsvp graceful
% Warning : RSVP process will stop and needs to restart manually,
You may loose ospf configuration, if not saved
Proceed for graceful restart? (y/n):y
```

revert-timer

Use this command to specify a "revert-timer" for the RSVP daemon. This command can be invoked from either the router rsvp mode or from the rsvp-trunk mode. In the RSVP router mode, this command defines the revert timer from secondary session to primary session for the RSVP daemon.

Use the `no` parameter with this command to revert to the default settings.

Command Syntax

```
revert-timer <1-65535>
no revert-timer (<1-65535>|)
```

Parameters

<1-65535>

Hold-on timer before revert back to primary session

Default

By default, `no revert-timer` is configured

Command Mode

Router or Trunk mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#revert-timer 30

#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#no revert-timer

#configure terminal
(config)#router rsvp
(config-router)#revert-timer 40
#configure terminal
(config)#router rsvp
(config-router)#no revert-timer
```

router rsvp

Use this command to enter router mode from configure mode and to enable the RSVP daemon, if it is not already enabled.

Use the `no` parameter with this command to disable RSVP on the node.

Command Syntax

```
router rsvp
no router rsvp
```

Parameters

None

Default

RSVP is started only if this command is executed.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The command prompt changes from config to config-router, as illustrated below:

```
#configure terminal
(config)#router rsvp
(config-router)#

(config-router)#exit
(config)#no router rsvp
```

rsvp hello-interval

Use this command to enable the sending of Hello packets on the interface and to set the interval value between successive Hello packets to neighbor. For optimum performance, set this value to less than one-third the value of the configured RSVP hello-timeout. See the [rsvp hello-timeout \(page 947\)](#) command for more information.



Note: This is an interface-specific command and when not used, the global hello-interval state applies.

Use the `no` parameter with this command to return to the default hello interval value.

Command Syntax

```
rsvp hello-interval <1-65535>
no rsvp hello-interval
```

Parameter

<1-65535>

RSVP hello interval in seconds

Default

By default, RSVP hello interval is 2 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#rsvp hello-interval 110

(config)#interface eth0
(config-if)#no rsvp hello-interval
```

rsvp hello-receipt

Use this command to enable the receipt of hello messages from peers connected through this interface. This is an interface-specific command and when not used, the global [hello-receipt \(page 899\)](#) command applies.

Use the `no` parameter with this command to disable the exchange of hello messages for this interface.

Command Syntax

```
rsvp hello-receipt  
no rsvp hello-receipt
```

Parameters

None

Default

By default, rsvp hello receipt is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#interface eth0  
(config-if)#rsvp hello-receipt
```

rsvp hello-timeout

This command determines how long an RSVP node should wait for a hello message before declaring a neighbor to be down. If an LSR does not received a hello message from a peer connected to an interface within the specified duration, the LSR resets all sessions that are shared with this particular peer. This is an interface-specific command and when not used, the global [hello-timeout \(page 900\)](#) command applies.

Use the `no` parameter to revert to the default hello timeout value.

Command Syntax

```
rsvp hello-timeout <1-65535>
no rsvp hello-timeout
```

Parameters

<1-65535>

Time to receive a hello message, in seconds

Default

By default, hello-timeout value is 7 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#rsvp hello-timeout 550

(config)#interface eth0
(config-if)#no rsvp hello-timeout
```

rsvp keep-multiplier

This command sets the constant for calculating a valid reservation lifetime for an LSP, which allows messages to be exchanged through this interface. This is an interface-specific command and when not specified, the global keep-multiplier command applies.

Reservation lifetime is the duration of bandwidth reservation for the LSP. The refresh time and keep multiplier are two interrelated timing parameters used to calculate the valid reservation lifetime for an LSP. Use the following formula to calculate the reservation lifetime for an LSP:

$$L \geq (K + 0.5) * 1.5 * R$$

K = keep-multiplier

R = refresh timer

Refresh messages are sent periodically so that neighbors do not timeout.

Use the `no` parameter with this command to return to the global keep-multiplier value.

Command Syntax

```
rsvp keep-multiplier <1-255>
no rsvp keep-multiplier <1-255>
```

Parameters

<1-255>

Set a value for the lifetime constant

Default

By default RSVP keep-multiplier value is 3.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#rsvp keep-multiplier 3

(config)#interface eth0
(config-if)#no rsvp keep-multiplier
```

rsvp refresh-time

Use this command to configure RSVP refresh interval timer for the current interface. This is an interface-specific command and when not used, the global [refresh-time \(page 938\)](#) command applies.

The refresh time and keep multiplier are two interrelated timing parameters used to calculate the valid reservation lifetime for an LSP. Refresh time regulates the interval between refresh messages that include path and reservation request (Resv) messages. Refresh messages are sent periodically so that the reservation does not timeout in the neighboring nodes. Each sender and receiver host sends path and resv messages, downstream and upstream respectively, along the paths.

Use the `no` parameter with this command to revert to the refresh-time value set in RSVP mode.

Command Syntax

```
rsvp refresh-time <1-65535>
no rsvp refresh-time <1-65535>
```

Parameter

<1-65535>

The duration for which messages are sent, in seconds

Default

By default, refresh interval is 30 seconds.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#rsvp refresh-time 5055

(config)#interface eth0
(config-if)#no rsvp refresh-time
```

rsvp-path

Use this command to create a new RSVP path or to enter the `Path` command mode. In this mode, you can add or delete paths and also specify the path to be loose or strict.

Use the `no` parameter with this command to delete the path and its specified hops.

Command Syntax

```
rsvp-path PATHNAME  
no rsvp-path PATHNAME
```

Parameter

PATHNAME

Name of the path

Default

By default, rsvp path is disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#rsvp-path mypath  
(config-path)#
```

rsvp-trunk

Use this command to create a new RSVP trunk. When the trunk is created, the attributes required to configure an explicitly-routed or traditionally-routed LSP are set. Once a trunk is configured with the required attributes, an RSVP session (and PSB) is created for this trunk, which enables the exchange of messages and completes the LSP setup.

This command also modifies an existing RSVP path to configure an explicitly-routed or traditionally-routed LSP. In addition, this command can be used to set the address family (IPv4) of an RSVP trunk. If no address family is assigned, the default value is used. If the address family is already set, a check is made to see whether the address family configured and the one already in the database are the same. An error message is returned if the two do not match.

Use the `no` parameter with this command to remove an RSVP trunk and all configured attributes, except the primary path.



Note: The RSVP trunk's name (TRUNKNAME) is limited to 32 characters.

Command Syntax

```
rsvp-trunk TRUNKNAME ipv4
no rsvp-trunk TRUNKNAME
```

Parameters

TRUNKNAME

Name to use for the trunk

ipv4

IPv4 address family trunk

Default

By default, rsvp trunk is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The command prompt changes from `config` to `config-trunk` as illustrated below:

```
#configure terminal
(config)#rsvp-trunk mytrunk ipv4
(config-trunk)#
```

rsvp-session force-reoptimize

Use this command to force re-optimize a specific RSVP session by providing its session name for any session (primary or secondary) within the trunk.

Command Syntax

```
rsvp-session SESSIONNAME force-reoptimize
```

Parameters

SESSIONNAME

Specify the name of the RSVP session to reoptimize. Must reference an existing RSVP session configured on the router. The session name is derived from the trunk name by appending suffixes such as “-Primary, -Secondary, -Detour”, etc.

force-reoptimize

Triggers re-optimization of the specified session.

Default

None

Command Mode

Execution mode

Applicability

Introduced before OcNOS version 1.3. In OcNOS version 7.0.0, the force-reoptimize option was updated to accept a session name as input and to re-optimize the specified session, supporting primary or secondary tunnels.

Examples

This example forces reoptimization of the secondary RSVP session `t1-Secondary`; the `Type` field in the `show rsvp session` output displays both the primary (PRI) and secondary (SEC) sessions.

```
#rsvp-session t1-Secondary force-reoptimize
```

```
OcNOS#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To      From      Tun-ID  LSP-ID  Type  LSPName  State
Uptime  Rt  Style  Labelin Labelout
3.3.3.3      2.2.2.2      5001    2203    PRI   t1-      24322
Primary
3.3.3.3      2.2.2.2      5001    2204    SEC   t1-      24321
Secondary
Total 2 displayed, Up 2, Down 0.
```

rsvp-trunk force-switchover-secondary

Use this command to force switchover a secondary LSP. This command is recommended for limited use on a stable system when there is a need for software upgrade without traffic impact.



Note: This command is supported on secondary configured tunnels and not on multiple secondary configured tunnels.

Command Syntax

```
rsvp-trunk TRUNKNAME force-switchover-secondary (off|on)
```

Parameters

TRUNKNAME

Name of the trunk

force-switchover-secondary

(off|on)

Force switchover to secondary

Default

None

Command Mode

Execute mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#rsvp-trunk <Trunk-Name> force-switchover-secondary off
```

rsvp-trunk-restart

Use this command to restart the RSVP trunk. This command “kills” an existing LSP and restarts the LSP setup process.

Command Syntax

```
rsvp-trunk-restart
```

Parameters

None

Default

By default, rsvp trunk restart is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#rsvp-trunk-restart
```

secondary

Use this command to enable secondary RSVP session.

Use the *no* parameter with this command to remove the secondary RSVP session.

.

Command Syntax

```
secondary  
no secondary
```

Parameters

None

Default

By default, Secondary is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

Example to delete secondary tunnel.

```
(config)#rsvp-trunk mytrunk  
(config-trunk)#secondary  
(config-trunk)#no secondary
```

secondary ADMIN-GROUP-NAME

Use this command to configure secondary administrative groups. Administrative groups are manually assigned attributes that describe the color of links, so that links with the same color are in one class. These groups are used to implement different policy-based LSP setups. Administrative group attributes can be included or excluded for an LSP or for a path's primary and secondary paths.



Note: A link can be added to a specific Administrative Group via NSM. Refer to the Network Services Module Command Reference for details.

Use the `no` parameter to remove a previously set group from an administrative group list.

Command Syntax

```
secondary (include-any|include-all|exclude-any) ADMIN-GROUP-NAME
secondary (include-any|exclude-any) ADMIN-GROUP-NAME
no secondary (include-any|include-all|exclude-any) ADMIN-GROUP-NAME
no secondary (include-any|exclude-any) ADMIN-GROUP-NAME
```

Parameters

include-any

Include any attribute

include-all

Include all attribute

exclude-any

Exclude any attribute

ADMIN-GROUP-NAME

Administrative group name

Default

By default, secondary admin group name is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#secondary exclude-any myadmingroup

#configure terminal
(config)#rsvp-trunk mytrunk
```



```
(config-trunk)#secondary include-any myadmingroup

#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#secondary include-all myadmingroup
```

secondary bandwidth

Use this command to reserve the bandwidth in bits per second for the current trunk.

Each LSP has an associated bandwidth attribute. The bandwidth value is included in the sender's RSVP Path message and specifies the bandwidth to be reserved for the LSP. It is set in bits per second, with a higher value indicating a greater user traffic volume. A zero bandwidth reserves no resources, although label exchanges are possible.

Use the `no` parameter with this command to unset the configured bandwidth information.

Command Syntax

```
secondary bandwidth BANDWIDTH  
no secondary bandwidth BANDWIDTH  
no secondary bandwidth
```

Parameter

BANDWIDTH

<1-999>k for 1 to 999 kilobits/s

<1-999>m for 1 to 999 megabits/s

<1-100>g for 1 to 100 gigabits/s

Default

By default, bandwidth is 0 bits per second, which allows data to flow through but does not reserve bandwidth.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#secondary bandwidth 100m
```

secondary cspf

Use this command to enable the use of Constrained Shortest Path First (CSPF) server for an explicit route to the egress, or all RSVP sessions.

The CSPF server computes paths for LSPs that are subject to constraints such as bandwidth, hop count, administrative groups, priority, and explicit routes. When computing paths for LSPs, CSPF considers not only the topology of the network and the attributes defined for the LSP, but also the links. It attempts to minimize congestion by intelligently balancing the network load.

Use the [no secondary cspf \(page 967\)](#) command to revert to the default settings.



Note: CSPF server information is not signaled across session and hence sessions in transit and egress nodes will not be aware of the CSPF server. So, in multi CSPF scenarios, neighbor down event from a CSPF server restart all sessions irrespective of which CSPF server sessions were using.

Command Syntax

```
secondary cspf
```

Parameters

None

Default

By default, secondary cspf is enabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows using the `no cspf` command in Trunk mode to disable CSPF for the primary LSP.

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#secondary cspf
```

secondary cspf-retry-limit

Use this command to specify the number of retries that CSPF should carry out for a request received from RSVP. Use the `no` parameter with this command to remove this configuration.

Command Syntax

```
secondary cspf-retry-limit <1-65535>  
no secondary cspf-retry-limit
```

Parameter

<1-65535>

The number of times CSPF should retry for this LSP

Default

By default, no retry limit for CSPF route calculations is configured, so the value is 0.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk T1  
(config-trunk)#secondary cspf-retry-limit 535
```

secondary cspf-retry-timer

Use this command to specify the time between each retry that CSPF might carry out for a request received from RSVP.

Use the `no` parameter with this command to remove this configuration.

Command Syntax

```
secondary cspf-retry-timer <1-600>  
no secondary cspf-retry-timer
```

Parameters

<1-600>

Timeout between successive retries, in seconds

Default

By default, no retry-timer configuration is defined for CSPF calculations, so the value is set to 0.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk T1  
(config-trunk)#secondary cspf-retry-timer 45
```

secondary filter

Use this command to set the filter to fixed or shared filter style for RSVP trunk.

- The shared filter style identifies a shared reservation environment. It creates a single reservation into which flows from all senders are mixed.
- The fixed filter style designates a distinct reservation. A distinct reservation request is created for data packets from a particular sender. The fixed filter style is also used style to prevent rerouting of an LSP and to prevent another LSP from using this bandwidth.

Use the `no` parameter to reset the configured filter to the default style.

Command Syntax

```
secondary filter fixed  
no secondary filter
```

Parameters

fixed

Use a Fixed Filter for this RSVP Trunk.

Default

By default, secondary filter is shared-explicit.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Usage

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#secondary filter fixed
```

secondary hold-priority

Use this command to configure the hold priority value for the selected trunk.

In case of insufficient bandwidth, the user must remove any less important existing LSP to free up the bandwidth. This can be done by preempting one or more of the signaled LSPs. Hold priority determines the degree to which an LSP holds onto its reservation for a session after the LSP has been configured successfully. When the hold priority is high, the existing LSP is less likely to give up its reservation.

Use the `no` parameter to revert to the default hold-priority value.

Command Syntax

```
secondary hold-priority <0-7>  
no secondary hold-priority
```

Parameter

<0-7>

Specify a value for hold priority

Default

The default hold-priority is 0, the highest value. Once a session is configured with a 0 hold priority value, no other session can preempt it.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#secondary hold-priority 2
```

secondary hop-limit

Use this command to specify a limit of hops for an RSVP trunk.

Upon configuration of an arbitrary hop-limit, the hop-limit is compared with the number of hops configured in the primary path, if a primary path has been configured. If the number of hops in the primary path exceed the hop-limit configured, no path messages are sent out and any existing session is torn down. If no primary path is configured, the trunk is processed normally and the path messages are sent out. The hop-limit data is sent to the CSPF server, if CSPF is being used.

Use the `no` parameter to revert to the default hop-limit value.

Command Syntax

```
secondary hop-limit <1-255>
no secondary hop-limit <1-255>
no secondary hop-limit
```

Parameter

<1-255>

The number of acceptable hops

Default

By default, secondary hop limit is 255

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#secondary hop-limit 23
```


secondary label-record

Use this command to record all labels exchanged between RSVP enabled routers during the reservation setup process. This command records all labels exchanged for an LSP from the ingress to the egress, and helps with debugging.

Use the `no` parameter to turn off recording.

Command Syntax

```
secondary label-record  
no secondary label-record
```

Default

By default, secondary label record is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#secondary label-record
```

no secondary affinity

Use this command to disable the use of sending out session attribute objects with resource affinity data.

Use the [secondary bandwidth \(page 958\)](#) command to revert to the default settings.

Command Syntax

```
no secondary affinity
```

Parameters

None

Default

By default, no secondary affinity is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#no secondary affinity
```

no secondary cspf

Use this command to disable the use of Constrained Shortest Path First (CSPF) server for an explicit route to the egress, or all RSVP sessions.

If CSPF is turned off globally, it cannot be enabled for any LSP. If used per LSP, it can be used to turn off CSPF computation for a specific LSP. The CSPF server computes paths for LSPs that are subject to various constraints such as bandwidth, hop count, administrative groups, priority, and explicit routes. When computing paths for LSPs, CSPF considers not only the topology of the network and the attributes defined for the LSP, but, also the links. It attempts to minimize congestion by intelligently balancing the network load.

Disable CSPF when all nodes do not support the required traffic engineering extensions and configure LSPs manually to use an explicit path. The LSP is then established only along the path specified by the operator.

Use the [secondary cspf \(page 959\)](#) command to revert to the default settings.

Command Syntax

```
no secondary cspf
```

Parameters

None

Default

By default, secondary no cspf is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows using the `no cspf` command in Trunk mode to disable CSPF for the primary LSP.

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#no secondary cspf
```

no secondary record

This command is used to disable recording of the route taken by path and resv messages and confirms the establishment of reservations and to identify errors. Routes are recorded by means of the route record object (RRO) in an RSVP message.

Use the [secondary record \(page 976\)](#) command to revert to the default settings.

Command Syntax

```
no secondary record
```

Parameters

None

Default

By default, routes are recorded

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#no secondary record
```

secondary path

Use this command to specify an RSVP path to be used.

Use the `no` parameter with this command to remove a configured RSVP path.



Notes:

- When a secondary path is configured (strict or loose), it is computed entirely based on the specified path constraints and does not consider disjointness from the primary path.
- If no explicit path is configured for the secondary, the system attempts to compute a fully disjoint path from the primary. The secondary path will not be established if such a disjoint path is unavailable.

Command Syntax

```
secondary path PATHNAME
no secondary path
```

Parameters

PATHNAME

The name of the path to be used. `PATHNAME` is a string (name) used to identify an RSVP path defined for the node (refer to the [rsvp-path \(page 950\)](#) command).

Default

By default, secondary path is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#secondary path mypath
```

secondary-priority path

Use this command to specify a RSVP path to be used for a specific priority secondary.

Use the no parameter with this command to remove a configured secondary-priority.

Command Syntax

```
secondary-priority <1-5> path PATHNAME  
no secondary-priority <1-5>
```

Parameters

<1-5>

Secondary Priority value.

PATHNAME

The name of the path to be used. **PATHNAME** is a string (name) used to identify an RSVP path defined for the node (refer to the [rsvp-path \(page 950\)](#) command).

Default

Secondary-priority can only be configured along with a path. Other attributes can only be associated post this command.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk t1  
(config-trunk)#secondary-priority 1 path sp1  
(config)# rsvp-trunk t1  
(config-trunk)#no secondary-priority 1
```

secondary-priority hold-priority

Use this command to configure the hold priority value for the secondary-priority lsp.

In case of insufficient bandwidth, the user must remove any less important existing LSP to free up the bandwidth. This can be done by preempting one or more of the signaled LSPs. Hold priority determines the degree to which an LSP holds onto its reservation for a session after the LSP has been configured successfully. When the hold priority is high, the existing LSP is less likely to give up its reservation.

Use the no parameter to revert to the default hold-priority value.

Command Syntax

```
secondary-priority <1-5> hold-priority <0-7>  
no secondary-priority <1-5> hold-priority <0-7>  
no secondary-priority <1-5> hold-priority
```

Parameters

<1-5>

Secondary Priority value.

<0-7>

Specify a value for hold priority.

Default

The default hold-priority is 0, the highest value. Once a session is configured with a 0 hold priority value, no other session can preempt it.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)# rsvp-trunk t1  
(config-trunk)#secondary-priority 1 hold-priority 4  
(config)# rsvp-trunk t1  
(config-trunk)# no secondary-priority 1 hold-priority
```

secondary-priority setup-priority

Use this command to configure a setup priority value for the secondary-priority lsp.

In case of insufficient bandwidth, the user must remove any less important LSPs to free up bandwidth. This can be done by preempting one or more of the existing LSPs. The setup priority determines whether a new LSP that preempts an existing LSP may be established. The setup priority of the new LSP must be higher than the hold priority of an existing LSP for the existing LSP to be preempted. Note that for a trunk, the setup priority should not be higher than the

hold priority.

Use the no parameter with this command to revert to the default setup priority value.

Command Syntax

```
secondary-priority <1-5> setup-priority <0-7>  
no secondary-priority <1-5> setup-priority <0-7>
```

Parameters

<1-5>

Secondary Priority value.

<0-7>

Specify a value for hold priority.

Default

By default, setup value is 7 (the lowest).

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)# rsvp-trunk t1  
(config-trunk)#secondary-priority 1 setup-priority 4  
(config)# rsvp-trunk t1  
(config-trunk)# no secondary-priority 1 setup-priority 4
```


secondary-priority label-record

Use this command to record all labels exchanged between RSVP enabled routers during the reservation setup process. This command records all labels exchanged for an LSP from the ingress to the egress, and helps with debugging.

Use the no parameter to turn off recording.

Command Syntax

```
secondary-priority <1-5> label-record  
no secondary-priority <1-5> label-record
```

Parameters

<1-5>

Secondary Priority value.

Default

By default, label record is disabled for secondary-priority.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)# rsvp-trunk t1  
(config-trunk)#secondary-priority 1 label-record  
(config)# rsvp-trunk t1  
(config-trunk)# no secondary-priority 1 label-record
```

secondary-priority hop-limit

Use this command to specify a limit of hops for a secondary-priority lsp.

Upon configuration of an arbitrary hop-limit, the hop-limit is compared with the number of hops configured in the path, if a path has been configured. If the number of hops in the path exceed the hop-limit configured, no path messages are sent out and any existing session is torn down.

Use the no parameter to revert to the default hop-limit value.

Command Syntax

```
secondary-priority <1-5> hop-limit <1-255>  
no secondary-priority <1-5> hop-limit <1-255>
```

Parameters

<1-5>

Secondary Priority value.

<1-255>

The number of acceptable hops.

Default

By default, hop limit is 255.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)# rsvp-trunk t1  
(config-trunk)#secondary-priority 1 hop-limit 123  
(config)# rsvp-trunk t1  
(config-trunk)# no secondary-priority 1 hop-limit 123
```

secondary-priority bandwidth

Use this command to reserve the bandwidth in bits per second for the current trunk.

Each LSP has an associated bandwidth attribute. The bandwidth value is included in the sender's RSVP Path message and specifies the bandwidth to be reserved for the LSP. It is set in bits per second, with a higher value indicating a greater user traffic volume. A zero bandwidth reserves no resources, although label exchanges are possible.

Use the no parameter with this command to unset the configured bandwidth information.

Command Syntax

```
secondary-priority <1-5> bandwidth BANDWIDTH  
no secondary-priority <1-5> bandwidth BANDWIDTH
```

Parameters

<1-5>

Secondary Priority value.

BANDWIDTH

<1-999>k for 1 to 999 kilobits/s

<1-999>m for 1 to 999 megabits/s

<1-100>g for 1 to 100 gigabits/s

Default

By default, bandwidth is 0 bits per second, which allows data to flow through but does not reserve bandwidth.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
(config)# rsvp-trunk t1  
(config-trunk)#secondary-priority 1 bandwidth 100m  
(config)# rsvp-trunk t1  
(config-trunk)# no secondary-priority 1 bandwidth 100m
```

secondary record

This command is used to enable recording of the route taken by path and resv messages to confirm the establishment of reservations and to identify errors. Routes are recorded by means of the route record object (RRO) in RSVP messages.

Use the [no secondary record \(page 968\)](#) command to revert to the default settings.

Command Syntax

```
secondary record
```

Parameters

None

Default

By default, routes are recorded

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#secondary record
```

secondary retry-limit

Use this command to specify a retry count this RSVP Trunk.

If a session is in a nonexistent state due to the receipt of a path error message, it tries to recreate the LSP for the number of times specified by [primary retry-limit \(page 933\)](#). Although the same retry command controls both the trunk and the session, the retry-limit value affects only the session and not the trunk. If the trunk is in an incomplete state, the code keeps trying to bring it to a complete state, irrespective of the retry-limit value.

Use the `no` parameter to revert to the default retry-limit value.

Command Syntax

```
secondary retry-limit <1-65535>
no secondary retry-limit
```

Parameter

<1-65535>

The set number of times the system should try setting up the LSP

Default

By default, the retry-limit value is 0 so the trunk and session try to create the LSP indefinitely.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#secondary retry-limit 256
```

secondary retry-timer

Use this command to specify a retry interval for an RSVP Trunk. When the ingress tries to configure an LSP and the setup fails due to the receipt of a path error message, the system waits for the time configured by this command before retrying the LSP setup process.

Use the `no` parameter to revert to the default.

Command Syntax

```
secondary retry-timer <1-600>  
no secondary retry-timer
```

Parameter

<1-600>

Interval after which the system should retry setting up the LSP, in seconds

Default

By default, retry time is 30 seconds

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#secondary retry-timer 12
```

secondary reuse-route-record

Use this command to use the updated route record list as an explicit route (with all strict nodes) when a path message is sent out at the next refresh.

An explicit route object (ERO) list contains the hops to be taken to reach the egress from the current LSR. If CSPF can not place an ERO with all strict routes, then this command helps modify the ERO after receiving resv messages. Future path messages have the ERO with all strict nodes, which identify each and every node to be traversed.

Use the `no` parameter to disable the use of the route record list as the explicit route.

Command Syntax

```
secondary reuse-route-record  
no secondary reuse-route-record
```

Parameters

None

Default

By default, secondary reuse route record is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#secondary reuse-route-record
```

secondary setup-priority

Use this command to configure a setup priority value for this trunk.

In case of insufficient bandwidth, the user must remove any less important LSPs to free up bandwidth. This can be done by preempting one or more of the existing LSPs. The setup priority determines whether a new LSP that preempts an existing LSP may be established. The setup priority of the new LSP must be higher than the hold priority of an existing LSP for the existing LSP to be preempted. Note that for a trunk, the setup priority should not be higher than the hold priority.

Use the `no` parameter with this command to revert to the default setup priority value.

Command Syntax

```
secondary setup-priority <0-7>  
no secondary setup-priority
```

Parameters

<0-7>

The priority value

Default

By default, setup value is 7 (the lowest).

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk mytrunk  
(config-trunk)#secondary setup-priority 2
```


secondary traffic

Use this command to identify the traffic type for this RSVP Trunk.

Use the `no` parameter with this command to unset the configured traffic type.

Command Syntax

```
secondary traffic (guaranteed|controlled-load)
no secondary traffic
```

Parameters

guaranteed

Guaranteed traffic

controlled-load

Controlled load traffic

Default

Controlled load is the default traffic type.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#secondary traffic guaranteed
```

session state timeout

Use this command to configure the maximum time a session can wait in a few transitory states for the acknowledgment or event to change the state. By default, the timeout value is 5 minutes. If no acknowledgment or event received within this period, session will timeout and proceed with release.

Use the *no* parameter with this command to revert to the default settings.

.

Command Syntax

```
session-state-timeout <5-60>
no session-state-timeout
```

Parameters

<5-60>

Timeout value in minutes

Default

None.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 7.0.0.

Example

```
#configure terminal
(config)#router rsvp
(config-router)# session-state-timeout 10
(config-router)# commit
```

```
#configure terminal
(config)#router rsvp
(config-router)# no session-state-timeout
(config-router)# commit
```

shutdown

Use this command to administratively disable a trunk. Execute this command to release the sessions of this trunk, and it only comes up once it is administratively enabled.

Use the `no` parameter to enable the trunk.

Command Syntax

```
shutdown
no shutdown
```

Parameters

None

Default

None

Command Mode

Trunk mode

Applicability

Introduced in OcNOS version 6.5.1.

Examples

The below example disable the trunk administratively:

```
#configure terminal
(config)#rsvp-trunk t1
(config-trunk)#shutdown
(config-trunk)#commit
```

The below example enable the trunk administratively:

```
#configure terminal
(config)#rsvp-trunk t1
(config-trunk)#no shutdown
(config-trunk)#commit
```

snmp restart rsvp

Use this command to restart SNMP in Resource Reservation Protocol -Traffic Engineering (RSVP-TE)

Command Syntax

```
snmp restart rsvp
```

Parameters

None

Default

By default, snmp restart rsvp is disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#snmp restart rsvp
```

to A.B.C.D

Use this command to specify an IPv4 egress for an LSP. When configuring an LSP, you must specify the address of the egress router by using this command in the trunk node. An egress definition is a mandatory attribute; no RSVP session is created when an egress is not defined.

Use the `no` parameter with this command to unset the configured egress address.

Command Syntax

```
to A.B.C.D no to
```

Parameters

None

Default

The operator must specify an egress for LSP initialization to begin.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk mytrunk
(config-trunk)#to 10.10.0.5
```

update-type

Use this command to change the method of updating attributes for sessions (primary/ secondary) for this trunk.

- If make-before-break is configured (default type), a new LSP is created for each attribute update. When the new LSP becomes operational, the original LSP is torn down.
- If break-before-make is configured, the existing LSP is torn down and restarted for each attribute update.

Use the `no` parameter with this command to remove an update type.

Command Syntax

```
update-type (make-before-break|break-before-make)
no update-type
```

Parameters

make-before-break

Make before break update

break-before-make

Break before make update

Default

By default, make-before-break types of updates are carried out.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#update-type break-before-make

#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#update-type make-before-break
```

Fast Reroute Commands

This chapter describes the RSVP-TE Fast Reroute commands.



Note: Constrained Shortest Path First (CSPF) is mandatory for FRR to work.

| | |
|---|-----|
| detour-allow-primary-upstream-path | 988 |
| default-frr-protection | 989 |
| detour-identification | 990 |
| primary fast-reroute admin-group-name | 991 |
| primary fast-reroute bandwidth | 992 |
| primary fast-reroute hold-priority | 993 |
| primary fast-reroute hop-limit | 994 |
| primary fast-reroute node-protection | 995 |
| primary fast-reroute protection | 996 |
| primary fast-reroute setup-priority | 997 |

detour-allow-primary-upstream-path

Use this command to ensure detour formation to consider the upstream path of protected LSPs. This is a deviation to RFC 4090 section 6.2 recommendation (<https://datatracker.ietf.org/doc/html/rfc4090>). This command is intended to be used in special cases where detour protection is required on ring topology if no alternate path is available.

Use the no parameter with this command to bypass the upstream path to the protected LSP when choosing a detour path.



Note: This command is intended to be used in ring topology if detour support is required at the cost of resource and link bandwidth. This command is not recommended to be configured otherwise.

Command Syntax

```
detour-allow-primary-upstream-path
no detour-allow-primary-upstream-path
```

Parameters

None

Default

By default, detour formation excludes the protected LSP upstream path as per RFC 4090 section 6.2 recommendations.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.4.1

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#detour-allow-primary-upstream-path
(config-router)#commit
(config-router)#no detour-allow-primary-upstream-path
(config-router)#commit
```


default-frr-protection

Use this command to configure the default method of fast reroute protection when sender has not specified a method via FRR object but asked for local protection. This command is particularly useful with interop with Cisco as Cisco doesn't send FRR object in path message. By default, default FRR protection considered to be one-to-one in OcNOS and in case of interop with Cisco where default protection needed is facility, this command shall be configured on all OcNOS devices in the network.



Notes:

- Having this command configured in one OcNOS device and not configured in other OcNOS device in the network will cause unpredictable behavior as RFC recommendation for merge node behavior of facility and one-to-one are different.
- This command is applicable only when path message contains local protection flag set but doesn't contain FRR object. When FRR object mentions protection type explicitly, this command is not applicable and also, if path message doesn't request local protection, then also this command is not applicable.

Command Syntax

```
default-frr-protection (one-to-one | facility)
no default-frr-protection
```

Parameters

facility

Facility Backup (Bypass) protection

one-to-one

One-to-One protection mechanism

Default

By default, if local protection requested but FRR object not available, one-to-one protection is considered.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.3.1.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)# default-frr-protection facility
(config-router)# commit
(config-router)# no default-frr-protection
(config-router)# commit
```

detour-identification

Use this command to set a path-specific detour LSP identification method, using the detour object.

Use the no parameter with this command to unset the detour LSP identification method.



Note: This command helps identify the backup LSP identification method for one-to-one protection only.

Command Syntax

```
detour-identification (path|sender-template)
no detour-identification
```

Parameters

path

Set a path-specific detour identification method

sender-template

Set a sender template-specific detour identification method

Default

By default, detour identification is sender template

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#detour-identification path

#configure terminal
(config)#router rsvp
(config-router)#detour-identification sender-template

#configure terminal
(config)#router rsvp
(config-router)#no detour-identification
```

primary fast-reroute admin-group-name

Use this command to configure admin group preference on the backup path. This information is propagated through all the nodes of label switched path.



Note: Fast-reroute admin group is not supported for bypass. It is currently supported only for detours.

Use the `no` parameter to remove the admin group configuration for fr backup path.

Command Syntax

```
primary fast-reroute (include-any|exclude-any) ADMIN-GROUP-NAME
no primary fast-reroute (include-any|exclude-any) ADMIN-GROUP-NAME
```

Parameter

include-any

Include any attribute

exclude-any

exclude any attribute

ADMIN-GROUP-NAME

Administrative group name

Default

By default, fast reroute admin group name is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#primary fast-reroute include-any RED
```

primary fast-reroute bandwidth

Use this command to set the detour LSP bandwidth.



Note: This command helps identify attributes of the FRR backup LSP for the one-to-one protection method.

Use the `no` parameter with this command to unset fast-reroute LSP bandwidth.

Command Syntax

```
primary fast-reroute bandwidth BANDWIDTH
no primary fast-reroute bandwidth BANDWIDTH
no primary fast-reroute BANDWIDTH
```

Parameter

BANDWIDTH

<1-999>k for 1 to 999 kilobits/s

<1-999>m for 1 to 999 megabits/s

<1-100>g for 1 to 100 gigabits/s

Default

By default, primary fast reroute bandwidth is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#primary fast-reroute bandwidth 10000000
```

primary fast-reroute hold-priority

Use this command to set the hold-priority for a detour LSP.



Note: This command helps identify attributes of the FRR backup LSP for the one-to-one protection method.

Use the `no` parameter with this command to unset the detour LSP hold-priority.

Command Syntax

```
primary fast-reroute hold-priority <0-7>
no primary fast-reroute hold-priority (<0-7>|)
```

Parameter

<0-7>

Set the value for hold-priority

Default

By default, primary fast reroute hold priority is 0.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#primary fast-reroute hold-priority 3
```

primary fast-reroute hop-limit

Use this command to set the hop-limit for a detour LSP.



Note: This command helps identify attributes of the FRR backup LSP for the one-to-one protection method.

Use the `no` parameter with this command to unset the detour LSP hop-limit.

Command Syntax

```
primary fast-reroute hop-limit <1-255>  
no primary fast-reroute hop-limit (<1-255>|)
```

Parameter

<1-255>

Set the number of hops

Default

By default, primary fast reroute hop limit is 255.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk T1  
(config-trunk)#primary fast-reroute hop-limit 25
```

primary fast-reroute node-protection

Use this command to set node protection.



Note: This command helps identify attributes of the FRR backup LSP for the one-to-one protection method.

Use the `no` parameter with this command to remove node protection.

Command Syntax

```
primary fast-reroute node-protection
no primary fast-reroute node-protection
```

Parameters

None

Default

By default, primary fast reroute node protection is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#primary fast-reroute node-protection
```

primary fast-reroute protection

Use this command to create an Fast Reroute backup and to set an LSP one-to-one protection or facility backup mechanism.



Note: Traffic switching of less than 50ms is not applicable for unknown unicast traffic received on edge nodes for VPLS services.

Use the `no` parameter with this command to remove LSP protection mechanism.

Parameters

None

Command Syntax

```
primary fast-reroute protection (one-to-one|facility)
no primary fast-reroute protection (one-to-one|facility)
```

Parameters

one-to-one

One-to-one protection

facility

Facility backup (bypass) protection"

Default

By default, primary fast reroute protection is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#primary fast-reroute protection one-to-one
```

primary fast-reroute setup-priority

Use this command to configure a setup-priority for the detour LSP.



Note: This command helps identify attributes of the FRR backup LSP for the one-to-one protection method.

Use the `no` parameter with this command to remove the detour LSP setup-priority.

Command Syntax

```
primary fast-reroute setup-priority <0-7>  
no primary fast-reroute setup-priority (<0-7>|)
```

Parameter

<0-7>

Set a value for setup priority

Default

By default, primary fast reroute setup priority is 0.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk T1  
(config-trunk)#primary fast-reroute setup-priority 2
```

Refresh Reduction Commands

This chapter describes the RSVP-TE Refresh Reduction commands:

| | |
|------------------------------|------|
| ack-send-timer | 999 |
| ack-wait-timeout | 1000 |
| message-ack | 1001 |
| refresh-reduction | 1002 |
| rsvp ack-wait-timeout | 1003 |
| rsvp message-ack | 1004 |
| rsvp refresh-reduction | 1005 |

ack-send-timer

Use this command to configure the timer to send an acknowledgment message. Timer configuration increases the chances of piggy backing multiple acknowledgment messages but also adds delay in acknowledgment received by neighbor node. So, use this command with exact knowledge of optimum time.



Note: Configure this command with a value within the limit of ack-wait-timeout to avoid frequent timeout.

Command Syntax

```
ack-send-timer <1-5>
no ack-send-timer
```

Parameter

<1-5>

Value in seconds for acknowledgment send timer.

Default

By default, acknowledgment message is transmitted immediately without piggy-backing.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)# ack-send-timer 1
(config-router)# commit
(config-router)# no ack-send-timer
(config-router)# commit
```

ack-wait-timeout

Use this command to set the acknowledgement wait timeout for the RSVP daemon. This command can be invoked from config-router mode.

Use the `no` parameter with this command to revert to the default settings.

Command Syntax

```
ack-wait-timeout <1-65535>
no ack-wait-timeout
```

Parameters

<1-65535>

Configure acknowledgment wait timeout.

Default

By default, `no ack-wait-timeout` is enabled.

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 4.2.

Examples

```
#configure terminal
(config)router rsvp
(config-router)# ack-wait-timeout 30

#configure terminal
(config)router rsvp
(config-router)# no ack-wait-timeout
```

message-ack

Use this command to enable message acknowledgment for all messages being sent to neighbors that are known to support refresh reduction.

Use the `no` parameter with this command to disable message acknowledgment for all messages being sent to neighbors.

Command Syntax

```
message-ack
no message-ack
```

Parameters

None

Default

By default, message acknowledgment is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#message-ack

(config)#router rsvp
(config-router)#no message-ack
```

refresh-reduction

Use this command to enable refresh reduction capability advertisement for all interfaces.

Use the `no` parameter with this command disable refresh reduction capability advertisement for all interfaces.

Command Syntax

```
refresh-reduction
no refresh-reduction
```

Parameters

None

Default

By default, refresh reduction mechanism is enabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#refresh-reduction

(config)#router rsvp
(config-router)#no refresh-reduction
```

rsvp ack-wait-timeout

Use this command to configure the acknowledgment wait timeout for all neighbors detected via the specific interface.

Use the `no` parameter with this command to revert to the default acknowledgment wait timeout for the specified interface.

Command Syntax

```
rsvp ack-wait-timeout <1-65535>
no rsvp ack-wait-timeout
```

Parameters

<1-65535>

Specify a value for the acknowledgment wait timeout in seconds. The default timeout value is 10 seconds.

Default

By default, RSVP acknowledgment wait timeout is 10 seconds

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#rsvp ack-wait-timeout 5

(config)#interface eth0
(config-if)#no rsvp ack-wait-timeout
```

rsvp message-ack

Use this command to enable message acknowledgment for all messages being sent to the neighbors that have been detected via the specific interface.

Use the `no` parameter with this command to disable message acknowledgment for all messages being sent to the neighbors that have been detected via the specified interface.

Command Syntax

```
rsvp message-ack  
no rsvp message-ack
```

Parameters

None

Default

By default, message acknowledgment is disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#interface eth0  
(config-if)#rsvp message-ack  
  
(config)#interface eth0  
(config-if)#no rsvp message-ack
```


rsvp refresh-reduction

Use this command to enable Refresh Reduction capability advertisement for a specific interface.

Use the `no` parameter with this command disable Refresh Reduction capability advertisement for the specified interface.

Command Syntax

```
rsvp refresh-reduction
no rsvp refresh-reduction
```

Parameters

None

Default

By default, Refresh Reduction mechanism is enabled for all interfaces.

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#interface eth0
(config-if)#rsvp refresh-reduction

(config)#interface eth0
(config-if)#no rsvp refresh-reduction
```

Facility Backup Commands

This chapter describes the RSVP-TE bypass commands for facility backup protection.

| | |
|--------------------------------------|------|
| backup-bw-type | 1007 |
| bandwidth | 1008 |
| bypass-lsp-addr-query-interval | 1009 |
| cspf-retry-limit | 1010 |
| cspf-retry-timer | 1011 |
| filter | 1012 |
| hold-priority | 1013 |
| hop-limit | 1014 |
| label-record | 1015 |
| no record | 1016 |
| path | 1017 |
| preemption-type | 1018 |
| record | 1019 |
| retry-limit | 1020 |
| retry-timer | 1021 |
| reuse-route-record | 1022 |
| rsvp-bypass | 1023 |
| setup-priority | 1024 |
| shutdown | 1025 |
| to A.B.C.D | 1026 |
| traffic | 1027 |

backup-bw-type

Use this command to select the bypass trunk bandwidth support type.

Bypass trunks of dedicated bandwidth type will serve only bandwidth protections requested LSPs. The total bandwidth requirement of served LSPs will be less than or equal to the bandwidth configured on the bypass trunk. If an LSP with bandwidth protection and higher setup priority requests protection and bypass doesn't have sufficient bandwidth available, then LSPs with lower hold priority will be preempted to serve the LSP with higher setup priority.

Use the `no` parameter to remove configured backup bandwidth type.

Command Syntax

```
backup-bw-type (dedicated | best-effort)
no backup-bw-type
```

Parameters

dedicated

Dedicated backup bandwidth support

best-effort

Best effort backup bandwidth support

Default

best-effort

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#backup-bw-type dedicated
```

bandwidth

Use this command to reserve the bypass bandwidth in bits per second for the current trunk.

Each LSP has an associated bandwidth attribute. The bandwidth value is included in the sender's RSVP Path message and specifies the bandwidth to be reserved for the LSP. It is specified in bits per second, with a higher value indicating a greater user traffic volume. A zero bandwidth reserves no resources, although exchanges labels.

Use the `no` parameter to remove configured bandwidth information.

Command Syntax

```
bandwidth BANDWIDTH
no bandwidth BANDWIDTH
no bandwidth
```

Parameter

BANDWIDTH

<1-999>k for 1 to 999 kilobits/s

<1-999>m for 1 to 999 megabits/s

<1-100>g for 1 to 100 gigabits/s

Default

The default bandwidth is 0 bits per second, which allows data to flow through but does not reserve bandwidth.

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#bandwidth 100m
(config-bypass)#no bandwidth
```

bypass-lsp-addr-query-interval

Use this command to set the interval at which bypass trunk must query CSPF for LSP address. This mechanism ensures to update bypass trunk LSP addresses regularly so that, it can verify regularly if it can protect any LSP requesting protection.

Use the `no` parameter with this command to reset the interval to default value.



Note: Reducing interval to lower values may impact performance.

Command Syntax

```
bypass-lsp-addr-query-interval <10-60>  
no bypass-lsp-addr-query-interval
```

Parameter

<10-60>

Set interval of bypass trunk querying LSP address

Default

By default, interval is set to 60 seconds.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal  
(config)#router rsvp  
(config-router)# bypass-lsp-addr-query-interval 50
```

cspf-retry-limit

Use this command to specify the number of retries that CSPF should carry out for a request received from RSVP. Use the `no` parameter with this command to disable this configuration.

Command Syntax

```
cspf-retry-limit <1-65535>
no cspf-retry-limit
```

Parameter

<1-65535>

Set the number of times CSPF should retry for this LSP

Default

By default, retry-limit is 0 which means infinite retry.

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#cspf-retry-limit 535

(config)#rsvp-bypass bp1
(config-bypass)#no cspf-retry-limit
```

cspf-retry-timer

Use this command to specify the time between each retry that CSPF might carry out for a request received from RSVP. Use the no parameter with this command to disable this configuration.

Command Syntax

```
primary cspf-retry-timer <1-600>  
no primary cspf-retry-timer
```

Parameter

<1-600>

Timeout between successive retries, in seconds

Default

By default, retry-timer is 0

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal  
(config)#rsvp-bypass bp1  
(config-bypass)#cspf-retry-timer 45  
  
(config)#rsvp-bypass bp1  
(config-bypass)#no cspf-retry-timer
```

filter

Use this command to set the filter to the fixed or shared style for an LSP.

- The shared filter style identifies a shared reservation environment. It creates a single reservation into which flows from all senders are mixed.
- The fixed filter style designates a distinct reservation. A distinct reservation request is created for data packets from a particular sender. The fixed filter style is also used style to prevent rerouting of an LSP and to prevent another LSP from using this bandwidth.

Use the `no` parameter to reset the configured filter to the default.

Command Syntax

```
filter fixed
no filter
```

Parameters

fixed

Use a fixed filter for this LSP

Default

By default, bypass filter is shared-explicit.

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#filter fixed
```

hold-priority

Use this command to configure the hold priority value for the selected bypass trunk. In case of insufficient bandwidth, remove less important existing LSPs to free up a portion of the bandwidth. This can be done by preempting one or more of the signaled LSPs. Hold priority determines the degree to which an LSP holds onto its reservation for a session after the LSP has been configured successfully. When the hold priority is high, the existing LSP is less likely to give up its reservation.

Use the `no` parameter to reset the trunk to the default hold-priority value.

Command Syntax

```
hold-priority <0-7>
no hold-priority
```

Parameters

<0-7>

Set a hold priority for the bypass LSP

Default

The default hold-priority value is 0, which is the highest. Once a session is configured with a hold priority of 0, no other session can preempt it.

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#hold-priority 2
```

hop-limit

Use this command to specify a limit of hops for an RSVP bypass trunk. Hop-limit data is sent to the CSPF server if CSPF is used.

Upon configuration of an arbitrary hop-limit, the hop-limit is compared with the number of hops configured in the bypass path, if a bypass path has been configured. If the number of hops in the bypass path exceeds the hop-limit configured, no Path messages are sent, and any existing session is torn down. If no bypass path is configured, the bypass trunk is processed normally and Path messages are sent.

Use the `no` parameter to reset the bypass trunk to the default hop-limit value.

Command Syntax

```
hop-limit <1-255>  
no hop-limit
```

Parameters

<1-255>

Set the number of acceptable hops for the LSP

Default

By default, bypass hop limit is 255

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0

Examples

```
#configure terminal  
(config)#rsvp-bypass bp1  
(config-bypass)#hop-limit 23
```

label-record

Use this command to record all labels exchanged between RSVP-enabled routers during the reservation setup process.

Use the `no` parameter with this command to turn off recording.

Command Syntax

```
label-record  
no label-record
```

Parameters

None

Default

By default, bypass label record is disabled

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal  
(config)#rsvp-bypass bp1  
(config-bypass)#label-record
```

no record

Use this command to disable recording of the route taken by Path and Reservation Request (Resv) messages that confirm establishment of reservations and are used to identify errors. The routes are recorded by means of the Route Record Object (RRO) in RSVP messages.

Command Syntax

```
no record
```

Parameters

None

Default

Routes are recorded by default.

Command Mode

RSVP Bypass mode

Examples

```
#configure terminal
(config)#rsvp-bypass bypassname
(config-bypass)#no record
```

path

Use this command to specify an RSVP path to be used. The PATHNAME in this command is the string (name) used to identify an RSVP path defined for the node (refer to rsvp-path command).

Use the `no` parameter with this command to remove a configured RSVP path.

Command Syntax

```
path PATHNAME
no path
```

Parameters

PATHNAME

The name of the path to use

Default

By default, bypass path is disabled

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#path mypath
```

preemption-type

Use this command to configure preemption type which decides the criteria to be considered in case of preemption. Use the `no` parameter to remove configured preemption type.

Command Syntax

```
preemption-type (less-lsp-preempted | less-unused-bandwidth)
no preemption-type
```

Parameters

less-lsp-preempted

Set preemption type to minimize number of LSPs preempted

less-unused-bandwidth

Set preemption type to ensure less bypass bandwidth unused

Default

By default, preemption type is set to less-lsp-preempted.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#preemption-type less-unused-bandwidth
```

record

Use this command to enable recording of the route taken by Path and Reservation Request (Resv) messages to confirm establishment of reservations and identify errors. Routes are recorded by means of the Route Record Object (RRO) in RSVP messages.

Use the `no record` command to disable recording of routes.

Command Syntax

```
record
```

Parameters

None

Default

By default, routes are recorded

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#record
```

retry-limit

Use this command to specify a retry count this RSVP bypass Trunk.

If a session is in a nonexistent state due to a path error message, the system tries to recreate the LSP for the number of times specified by the retry-limit command.

Although the same retry command controls both the trunk and the session, the retry-limit value affects only the session and not the trunk. If the trunk is in an incomplete state, the code keeps trying forever to bring it to a complete state regardless of the retry-limit value.

Use the `no` parameter with this command to revert to the default retry-limit value.

Command Syntax

```
retry-limit <1-65535>
no retry-limit
```

Parameter

<1-65535>

The set number of times the system should try setting up the LSP

Default

By default, the retry-limit value is 0, and the trunk and session try to create the LSP indefinitely.

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#retry-limit 256
```

retry-timer

Use this command to specify a retry interval for an RSVP bypass Trunk. When an ingress node tries to configure an LSP and the setup fails due to the receipt of a Path Error message, the system waits for the time configured with this command, before retrying the LSP setup process.

Use the `no` parameter with this command to revert to the default retry-time value.

Command Syntax

```
retry-timer <1-600>  
no retry-timer
```

Parameters

<1-600>

Time in seconds after which the system should retry setting up the LSP

Default

By default, retry-timer value is 30 seconds.

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal  
(config)#rsvp-bypass bp1  
(config-bypass)#retry-timer 12
```

reuse-route-record

Use this command to use the updated Route Record List as an Explicit Route (with all strict nodes) when a path message is sent out at the next refresh.

The ERO list contains the hops to be taken to reach the egress from the current LSR. If CSPF is not available, to place an ERO with all strict routes, use this command to modify the ERO after receiving the Resv message. The future Path messages have the ERO with all strict nodes, identifying each and every node to be traversed.

Use the `no` parameter with this command to disable the use of the Route Record List as the explicit route.

Command Syntax

```
reuse-route-record
no reuse-route-record
```

Parameters

None

Default

By default, reuse route record is disabled

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#reuse-route-record
```

rsvp-bypass

Use this command to create a new RSVP bypass trunk. When the bypass trunk is created, the attributes required to configure an explicitly-routed or traditionally-routed LSP are set. Once a trunk is configured with the required attributes, an RSVP bypass session (and PSB) is created for this trunk, which enables the exchange of messages and completes the LSP setup.

This command also modifies an existing RSVP path to configure an explicitly-routed or traditionally-routed LSP.

Use the `no` parameter with this command to remove an RSVP bypass trunk and all configured attributes.



Note: The RSVP bypass' name (BYPASSNAME) is limited to 32 characters.

Command Syntax

```
rsvp-bypass BYPASSNAME
no rsvp-bypass BYPASSNAME
```

Parameters

BYPASSNAME

Name to use for the bypass trunk

Default

By default, `rsvp bypass trunk` is disabled

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

The command prompt changes from `config` to `config-bypass` as illustrated below:

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#
```

setup-priority

Use this command to configure a setup priority value for a trunk. In case of insufficient bandwidth, users must remove less important LSPs to free up the bandwidth. This can be done by preempting one or more of the existing LSPs. The primary setup priority determines if a new LSP can preempt an existing LSP.

The setup priority of the new LSP must be higher than the hold priority of an existing LSP for the existing LSP to be preempted. Note that for a trunk, the setup priority should not be higher than the hold priority.

Use the `no` parameter with this command to revert to the default primary setup priority value.

Command Syntax

```
setup-priority <0-7>  
no setup-priority
```

Parameters

<0-7>

Set the priority value

Default

By default, setup priority is 7, which is the lowest.

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal  
(config)#rsvp-bypass bp1  
(config-bypass)#setup-priority 2
```

shutdown

Use this command to administratively disable a bypass trunk. Execute this command to release the bypass session of this trunk, and it only comes up once it is administratively enabled.

Use the `no` parameter to enable the trunk.

Command Syntax

```
shutdown
no shutdown
```

Parameters

None

Default

None

Command Mode

Bypass mode

Applicability

Introduced in OcNOS version 6.5.1.

Examples

The below example disable the trunk administratively:

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#shutdown
(config-bypass)#commit
```

The below example enable the trunk administratively:

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#no shutdown
(config-bypass)#commit
```

to A.B.C.D

Use this command to specify an IPv4 egress for a bypass LSP. When configuring an LSP, you must specify the address of the egress router by using this command in the bypass node. An egress definition is a mandatory attribute; no RSVP session is created when an egress is not defined.

Use the `no` parameter with this command to unset the configured egress address.

Command Syntax

```
to A.B.C.D
no to
```

Parameters

None

Default

The operator must specify an egress for LSP initialization to begin.

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#to 10.10.0.5
```

traffic

Use this command to specify the traffic type for this RSVP bypass Trunk.

Use the `no` parameter with this command to reset the configured traffic type.

Command Syntax

```
traffic (guaranteed|controlled-load)
no traffic
```

Parameters

controlled-load

Controlled loaded traffic

guaranteed

Guaranteed traffic

Default

By default, primary traffic type is controlled-load

Command Mode

Bypass mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#rsvp-bypass bp1
(config-bypass)#traffic guaranteed
```

Differentiated Services Commands

This chapter describes the RSVP Differentiated Services (DiffServ) commands.

| | |
|--------------------------------|------|
| map-route A.B.C.D | 1029 |
| map-route X:X::X:X | 1030 |
| override-diffserv | 1031 |
| primary class-to-exp-bit | 1032 |
| primary elsp-signaled | 1033 |
| primary llsp | 1034 |
| secondary elsp-signaled | 1035 |
| secondary llsp | 1036 |
| secondary map class | 1037 |
| show rsvp diffserv-info | 1038 |

map-route A.B.C.D

Use this command to map a IPv4 prefix route onto a trunk. This route is to be used for packets that are mapped to a specific RSVP trunk.

Use the `no` parameter with this command for unmapping routes from specified trunks.



Note: Do not configure the local address as a map route, as explicit validation is not done. The wrong configuration may impact traffic.

Command Syntax

```
map-route A.B.C.D A.B.C.D map-route A.B.C.D A.B.C.D CLASS map-route A.B.C.D/M map-route A.B.C.D/M  
CLASS no map-route A.B.C.D A.B.C.D no map-route A.B.C.D A.B.C.D CLASS no map-route A.B.C.D/M no map-  
route A.B.C.D/M CLASS
```

Parameters

A.B.C.D

Specify the IPV4 address to be mapped.

A.B.C.D

Specify a mask to be applied to the address being mapped.

A.B.C.D/M

Specify the IPV4 address to be mapped, with mask.

CLASS

Specify the DiffServ Class Name (for example, `be`, `ef` etc.) used for selecting incoming IP packets to be mapped to a specified RSVP trunk.

Default

By default, map route A.B.C.D is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal (config)#rsvp-trunk T1 (config-trunk)#map-route 1.1.2.2/24 be
```

map-route X:X::X:X

Use this command to map a IPv6 prefix route onto a trunk. This route is to be used for packets that are mapped to a specific RSVP trunk.

Use the `no` parameter with this command for unmapping routes from specified trunks.

Command Syntax

```
map-route X:X::X:X X:X::X:X
map-route X:X::X:X X:X::X:X CLASS
map-route X:X::X:X/M
map-route X:X::X:X/M CLASS
no map-route X:X::X:X X:X::X:X
no map-route X:X::X:X X:X::X:X CLASS
no map-route X:X::X:X/M
no map-route X:X::X:X/M CLASS
```

Parameters

X:X::X:X

Specify the IPv6 address to be mapped.

X:X::X:X

Specify a mask to be applied to the address being mapped.

X:X::X:X/M

Specify the IPv6 address to be mapped, with mask.

CLASS

Specify the DiffServ Class Name (for example, `be`, `ef` etc.) used for selecting incoming IPv6 packets to be mapped to a specified RSVP trunk.

Default

By default, map route X:X::X:X is disabled.

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#map-route 3ffe::/32 be
```

override-diffserv

Use this command to enable the Differentiated Services (Diff-Serv) override configuration.

If a Path message is received without a Diff-Serv object by a Diff-Serv enabled node, it can be interpreted either as a request for an E-LSP (EXP-Inferred-PSC LSP) or as a request for Non-Diff-Serv LSP. This command supports the override option and when configured, the LSR interprets a path message without a Diff-Serv object as a request for Non-Diff-Serv LSP.

Use the `no` parameter with this command disable this feature.

Command Syntax

```
override-diffserv
no override-diffserv
```

Parameters

None

Default

By default, override diffserv is disabled

Command Mode

Router mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router rsvp
(config-router)#override-diffserv
```

primary class-to-exp-bit

Use this command to configure a primary Per-Hop Behavior-Experimental (PHB-EXP) mapping to be used by an EXP-Inferred-PSC LSP (E-LSP). This mapping is different from the node level PHB-EXP mapping.

Use the `no` parameter with this command to remove a PHB-EXP mapping configuration from current E-LSP PHB-EXP mapping.

Command Syntax

```
primary class-to-exp-bit CLASS <0-7>  
no primary class-to-exp-bit CLASS <0-7>
```

Parameters

CLASS

Specify the DiffServ Class Name (for example, be, ef etc.) used for selecting incoming IP packets to be mapped to a specified RSVP trunk.

<0-7>

Exp bit which is to be mapped to this PHB.

Default

By default, primary map class is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk T1  
(config-trunk)#primary class-to-exp-bit af12 3  
  
(config)#rsvp-trunk T1  
(config-trunk)#no primary class-to-exp-bit af12 3
```

primary elsp-signaled

Use this command to configure a primary Differentiated Services (Diff-Serv (Diff-Serv) explicitly signaled EXP-Inferred-PSC LSP (E-LSP) interface.

The classes 1 to 7 are optional parameters that can be selected from node level Per-Hop Behavior (PHB-EXP) mapping as PHBs, which will then be used for an E-LSP. If you do not specify a class with this command, all classes will be selected for the E-LSP.

Use the no parameter with this command to remove the configuration.

Command Syntax

```
primary elsp-signaled
primary elsp-signaled CLASS1
primary elsp-signaled CLASS1 CLASS2
primary elsp-signaled CLASS1 CLASS2 CLASS3
primary elsp-signaled CLASS1 CLASS2 CLASS3 CLASS4
primary elsp-signaled CLASS1 CLASS2 CLASS3 CLASS4 CLASS5
primary elsp-signaled CLASS1 CLASS2 CLASS3 CLASS4 CLASS5 CLASS6
primary elsp-signaled CLASS1 CLASS2 CLASS3 CLASS4 CLASS5 CLASS6 CLASS7
no primary elsp-signaled
```

Parameter

CLASS<0-7>

Diffserv class alias. e.g.: be, ef, af11, etc.

Default

By default, primary elsp signaled is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#primary elsp-signaled cs2 cs5 cs6

(config)#rsvp-trunk T1
(config-trunk)#no primary elsp-signaled
```

primary llsp

Use this command to configure a primary Differentiated Services Label-Only-Inferred-PSC (Diff-Serv L-LSP) interface, which will use Diff-Serv Class as its PHB Scheduling Class (PSC).

Use the no parameter with this command to remove the Diff-Serv L-LSP configuration.

Command Syntax

```
primary llsp CLASS  
no primary llsp
```

Parameters

CLASS<0-7>

Diffserv class alias. e.g: be, ef, af11, etc.

Default

By default, primary llsp is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

This command is not available on Qumran (Q1 and Q2) series platforms.

Examples

```
#configure terminal  
(config)#rsvp-trunk T1  
(config-trunk)#primary llsp cs4  
  
(config)#rsvp-trunk T1  
(config-trunk)#no primary llsp
```

secondary elsp-signaled

Use this command to configure a secondary Diff-Serv (Differentiated Services) explicitly signaled E-LSP (EXP-Inferred-PSC LSP) interface. The classes 1 to 7 are optional parameters can be selected from the node level PHB-EXP (Per-Hop Behavior) mapping as PHBs. They will then be used for an E-LSP. If you do not specify a class with this command, all classes will be selected for the E-LSP.

Use the no parameter with this command to remove the configuration.

Command Syntax

```
secondary elsp-signaled
secondary elsp-signaled CLASS1
secondary elsp-signaled CLASS1 CLASS2
secondary elsp-signaled CLASS1 CLASS2 CLASS3
secondary elsp-signaled CLASS1 CLASS2 CLASS3 CLASS4
secondary elsp-signaled CLASS1 CLASS2 CLASS3 CLASS4 CLASS5
secondary elsp-signaled CLASS1 CLASS2 CLASS3 CLASS4 CLASS5 CLASS6
secondary elsp-signaled CLASS1 CLASS2 CLASS3 CLASS4 CLASS5 CLASS6 CLASS7
no secondary elsp-signaled
```

Parameters

CLASS<0-7>

Diffserv class alias. e.g: be, ef, af11, etc.

Default

By default, secondary elsp signaled is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#rsvp-trunk T1
(config-trunk)#secondary elsp-signaled class cs3 cs6 cs2 cs5

(config)#rsvp-trunk T1
(config-trunk)#no secondary elsp-signaled
```

secondary llsp

Use this command to configure a secondary Differentiated Services Label-Only-Inferred-PSC (Diff-Serv L-LSP) interface, which will use Diff-Serv Class as its PHB Scheduling Class (PSC).

Use the no parameter with this command to remove the Diff-Serv L-LSP configuration.

Command Syntax

```
secondary llsp CLASS  
no secondary llsp
```

Parameters

CLASS<0-7>

Diffserv class alias. e.g: be, ef, af11, etc.

Default

By default, secondary llsp is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

This command is not available on Qumran (Q1 and Q2) series platforms.

Example

```
#configure terminal  
(config)#rsvp-trunk T1  
(config-trunk)#secondary llsp class cs5  
  
(config)#rsvp-trunk T1  
(config-trunk)#no secondary llsp
```

secondary map class

Use this command to configure a secondary PHB-EXP (Per-Hop Behavior-Experimental) mapping to be used by an E-LSP (EXP-Inferred-PSC LSP). This mapping is different from the node level PHB-EXP mapping.

Use the no parameter with this command to remove a PHB-EXP mapping configuration from current E-LSP PHB-EXP mapping.

Command Syntax

```
secondary map class-to-exp-bit CLASS <0-7>  
no secondary map class-to-exp-bit CLASS <0-7>
```

Parameters

CLASS

Diff-Serv class (queue) mapped to the particular PHB.

Diffserv class alias e.g: be, ef, af11, etc.

<0-7>

Exp bit that is to be mapped to this PHB.

Default

By default, secondary map class is disabled

Command Mode

Trunk mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#rsvp-trunk T1  
(config-trunk)#secondary class-to-exp-bit cs4 3  
  
(config)#rsvp-trunk T1  
(config-trunk)#no secondary class-to-exp-bit cs4 3
```

show rsvp diffserv-info

Use this command to display node level Differentiated Services (Diff-Serv) configuration information. This information includes the node level PHB-EXP mapping configured for ELSP-signaled LSP.

Command Syntax

```
show rsvp diffserv-info
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

Following is a sample output of the `show rsvp diffserv-info` command.

```
#show rsvp diffserv-info
CLASS-EXP mapping:
CLASS    DSCP_value
c5  101000 0
be  000000 1
cs1 001000 2
cs3 011000 3
cs2 010000 4
cs4 100000 5
cs6 110000 6
cs7 111000 7
```

Here is the explanation of the "show command" output fields.

Table 33. show rsvp diffserv-info output fields

| Field | Description |
|-----------|--|
| CLASS | MPLS class type that corresponds to the DiffServ traffic engineering class. |
| EXP_value | Exp value is initialized at the ingress routing device only and overrides the rewrite configuration established for that forwarding class. |

RSVP Multipath Commands

This chapter is a reference for RSVP Multipath commands:

| | |
|--|------|
| description | 1040 |
| multipath-group | 1041 |
| rsvp-multipath | 1042 |
| to A.B.C.D | 1043 |
| show mpls counters rsvp multipath-name | 1044 |
| show mpls rsvp-multipath | 1045 |
| show rsvp multipath | 1046 |
| show running-config rsvp-multipath | 1047 |

description

Use this command to add a description to the multipath group or update an existing description.

Use the `no` parameter to remove the description.

Command Syntax

```
description LINE
no description
```

Parameter

LINE

Line describing the purpose of RSVP multipath Group

Default

By default, description is empty.

Command Mode

Multipath mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
(config)#rsvp-multipath mygroup
(config-multipath)# description "For example purpose"
(config-multipath)# no description
```

multipath-group

Use this command to configure a multipath group on the RSVP trunk. When multipath group is configured on the trunk, trunk will become member of multipath group.

Use the `no` parameter to remove multipath group from RSVP trunk.

Command Syntax

```
multipath-group GROUPNAME  
no multipath-group
```

Parameter

GROUPNAME

Name of the multipath group

Default

By default, parameter is not configured.

Command Mode

Trunk mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal  
(config)#rsvp-trunk to_dest1  
(config-trunk)# multipath-group mygroup  
(config-trunk)# no multipath-group
```

rsvp-multipath

Use this command to create a new multipath group or enter the existing group to update any group parameter.

Use the `no` parameter with this command to remove an RSVP multipath group.

Command Syntax

```
rsvp-multipath GROUPNAME  
no rsvp-multipath GROUPNAME
```

Parameter

GROUPNAME

Name of the multipath group

Default

By default, rsvp multipath is not configured.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

The command prompt changes from config to config-multipath as illustrated below:

```
#configure terminal  
(config)#rsvp-multipath mygroup  
(config-multipath)#exit  
(config)#no rsvp-multipath mygroup
```

to A.B.C.D

Use this command to specify an IPv4 egress to a multipath group. When configuring a multipath group, you must specify the address of the egress router by using this command. An egress definition is a mandatory attribute; multipath group is not allowed to be configured on RSVP trunks when an egress is not defined.

Use the no parameter to remove the IPv4 egress address from multipath group.

Command Syntax

```
to A.B.C.D
no to
```

Parameter

A.B.C.D

IPv4 address of multipath group destination

Default

By default, parameter is not configured.

Command Mode

Multipath mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
(config)#rsvp-multipath mygroup
(config-multipath)# to 3.3.3.3
(config-multipath)# no to
```

show mpls counters rsvp multipath-name

Use this command to get traffic statistics over member trunks of multipath group.



Notes:

- Qumran devices require hardware statistics profiles to be enabled explicitly for traffic statistics to be active. For tunnel statistics 'tunnel-lif' statistics profiles need to be enabled and system needs to be reloaded to take effect. Refer to the `hardware-profile statistics` command in the *OcNOS System Management Guide* for details.
- This statistic is not an explicit count of traffic over a multipath group. This statistic is collected over multipath members individually and any traffic individually carried by multipath members also counted under this output.

Command Syntax

```
show mpls counters rsvp multipath-name (NAME|)
```

Parameter

NAME

RSVP multipath group name

Command Mode

Exec mode and Privilege Exec mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#show mpls counters rsvp multipath-name test
Tunnel-id 5001 Extended Tunnel-ID 1.1.1.1 Egress 2.2.2.2
  lsp-name : t1-Primary                                [Ingress]
  lsp-ingress : 1.1.1.1                                lsp-id : 2201
  Rx pkts : n/a                                         Rx bytes : n/a
  Tx pkts : 3776248                                     Tx bytes : 5671925998

Tunnel-id 5002 Extended Tunnel-ID 1.1.1.1 Egress 2.2.2.2
  lsp-name : t2-Primary                                [Ingress]
  lsp-ingress : 1.1.1.1                                lsp-id : 2202
  Rx pkts : n/a                                         Rx bytes : n/a
  Tx pkts : 3776250                                     Tx bytes : 5671927500
```


show mpls rsvp-multipath

Use this command to display forwarder level information for a multipath group or all multipath groups.

Command Syntax

```
show mpls rsvp-multipath (NAME|)
```

Parameter

NAME

RSVP multipath group name

Command Mode

Exec mode and Privilege Exec mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
# show mpls rsvp-multipath
Codes: > - installed FTN, * - selected FTN, t - tunnel, R - RSVP-TE FTN

Multipath Name : mp1, ID : 101, Nhlfe Ix : 2
Active member count : 2, FEC : 2.2.2.2/32
Active member details :
-----
  Index  Code      FTN-ID  Nhlfe-ID  Tunnel-id  Pri  LSP-Type  Out-Label  Out-
Intf    ELC      Nexthop
  1      R(t)>    1       1         5002      Yes  LSP_
DEFAULT 24320    xe31     No        31.1.1.2
  2      R(t)>    2       3         5001      Yes  LSP_
DEFAULT 24321    xe11     No        11.1.1.2
```

show rsvp multipath

Use this command to display information for a multipath group or all multipath groups.

Command Syntax

```
show rsvp multipath (NAME|)
```

Parameter

NAME

RSVP multipath group name

Command Mode

Exec mode and Privilege Exec mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#show rsvp multipath

RSVP-multipath Name : mp1, ID : 101
Description : "multipath group from R1 to R2"
Member count : 3, Egress : 2.2.2.2/32
Member details :
-----
Trunk-ID   Trunk-name           Status
5001       to_R2_1               active
5002       to_R2_2               active
5003       to_R2_3               inactive

RSVP-multipath Name : mp2, ID : 102
Member count : 0, Egress : 3.3.3.3/32

RSVP-multipath Name : mp3, ID : 103
Member count : 0, Egress : N/A
```

show running-config RSVP-multipath

Use this command to check configuration details of multipath groups.

Command Syntax

```
show running-config rsvp-multipath (NAME|)
```

Parameter

NAME

Name of the multipath group

Command Mode

Exec mode and Privilege Exec mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#show running-config rsvp-multipath
!
rsvp-multipath mp1
  description "multipath group from R1 to R2"
  to 2.2.2.2
!
rsvp-multipath mp2
  to 3.3.3.3
!
rsvp-multipath mp3
!
!
```

RSVP-TE Show Commands

This chapter describes the RSVP-TE show commands.

| | |
|---|------|
| show rsvp label-pool | 1049 |
| show debugging rsvp | 1050 |
| show rsvp | 1051 |
| show rsvp admin-groups | 1053 |
| show rsvp bypass | 1054 |
| show rsvp bypass detail | 1055 |
| show rsvp bypass lsp-address-list | 1057 |
| show rsvp bypass protected-lsp-list | 1058 |
| show rsvp control-adjacency | 1060 |
| show rsvp data-link | 1062 |
| show rsvp graceful-restart | 1063 |
| show rsvp interface | 1064 |
| show rsvp l2-info | 1066 |
| show rsvp local-addresses | 1067 |
| show rsvp lsp-id-table | 1068 |
| show rsvp neighbor | 1070 |
| show rsvp nexthop-cache | 1072 |
| show rsvp path | 1073 |
| show rsvp protected-lsp-reop-list | 1074 |
| show rsvp session | 1075 |
| show rsvp session count | 1080 |
| show rsvp session egress | 1081 |
| show rsvp session ingress | 1085 |
| show rsvp session LSP-NAME | 1089 |
| show rsvp session transit | 1092 |
| show rsvp statistics | 1095 |
| show rsvp summary-refresh | 1096 |
| show rsvp trunk | 1097 |
| show rsvp trunk multi-sec-detail | 1099 |
| show rsvp trunk-id-table | 1100 |
| show rsvp version | 1102 |

show rsvp label-pool

Use this command to display the label management in RSVP protocol.

Command Syntax

```
show rsvp label-pool (block-id <0-1638>|)
```

Parameter

block-id

Block identifier of a block that is allocated to the protocol

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 6.5.4.

Example

The following is an output from the show rsvp label-pool command.

```
#show rsvp label-pool
Module: RSVP, Label range: 16-1048575, Current block: 45
+-----+-----+-----+-----+-----+-----+
block_id label_min label_max usable_labels free_labels first_free_label
+-----+-----+-----+-----+-----+-----+
44 28160 28799 640 0 0
45 28800 29439 640 96 29344
Total - blocks: 2, used-labels: 1184, free-labels: 96
```

show debugging rsvp

This command displays the status of the options selected by the `debug` RSVP command.

Command Syntax

```
show debugging rsvp
```

Parameters

None

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show debugging rsvp
NSM debugging status:
  RSVP event debugging is on
  RSVP packet debugging is on
  RSVP incoming packet debugging is on
  RSVP outgoing packet debugging is on
  RSVP hexadecimal dump debugging is on
#
```

[Table 34](#) explains the show command output fields.

Table 34. show debugging rsvp output fields

| Field | Description |
|----------------------|--|
| NSM debugging status | Debugging is enabled or disabled on a per-interface basis, using the commands. |

show rsvp

Use this command to display data about the RSVP daemon.

Command Syntax

```
show rsvp
```

Parameters

None

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp
RSVP Version           : 1
Process uptime         : 8 minutes
RSVP Refresh Reduction : Enabled
RSVP Message Acknowledgement : Disabled
Bundle Send           : Disabled
NSM Connection         : Up
CSPF Connection        : Up
CSPF usage             : Enabled
RSVP Refresh Timer     : 5
Keep Multiplier        : 3
Acknowledgement Await Timeout : 10
Explicit-Null For Direct Conn : Disabled
Local Protection       : Disabled
Hello Receipt          : Disabled
Hello Interval         : 2
Hello Timeout          : 10
Loop detection         : Enabled (all interface)
Override Diffserv      : Disabled
Ingress                : 1.1.1.1
Penultimate Hop Popping : Enabled
Refresh PATH msg parsing : Enabled
Refresh RESV msg parsing : Enabled
Detour identification   : Sender-Template

#
```

[Table 35](#) explains the show command output fields.

Table 35. show rsvp output fields

| Field | Description |
|--------------|--|
| RSVP Version | Version number associated with the RSVP ingress route. |

Table 35. show rsvp output fields (continued)

| Field | Description |
|-------------------------------|--|
| Process uptime | Duration of the process running time. |
| RSVP Refresh Reduction | Measure of processing over head requests of refresh messages. Refresh reduction detail extensions improve routing device performance by reducing the process overhead, thus increasing the number of LSPs a routing device can support. |
| RSVP Message Acknowledgement | Acknowledge message for refresh reductions. |
| Bundle Send | Disables sending of Bundle Messages for a system. |
| NSM Connection | The Network Services Module (NSM) sends unsolicited messages to, or receives unsolicited messages from, the QoS (quality of service) module. |
| CSPF Connection | NSM passes the information to CSPF. |
| CSPF usage | CSPF finds the shortest path toward the LSP's egress router, taking into account explicit-path constraints. |
| RSVP Refresh Timer | Time interval used to generate periodic RSVP messages. |
| Keep Multiplier | Number of RSVP messages that can be lost before an RSVP state is declared stale. |
| Acknowledgment Await Timeout | The router that initiates the acknowledgment messages for an RSVP session waits for the timeout. |
| Explicit-Null For Direct Conn | Advertise label 0 to the egress routing device of an LSP. Explicit null: enabled or disabled. |
| Local Protection | A local repair mechanism is in use to maintain this tunnel. |
| Hello Receipt | To exchange Hello messages among neighbors. |
| Hello Interval | Frequency at which RSVP hellos are sent on this interface (in seconds). |
| Hello Timeout | RSVP Hello State Timer feature detects when a neighbor is down and triggers faster state timeout. |
| Loop detection | Loop back Detection (LBD) provides protection against loops by transmitting loop protocol packets out of ports where loop protection has been enabled. |
| Override Diffserv | Diffserv helps to carry the EXP-to-PHB mapping for signaled E-LSP or the PSC value for L-LSP. |
| Ingress | Information about ingress RSVP sessions. |
| Penultimate Hop Popping | Removes the label one hop before its destination. |
| Refresh PATH msg parsing | Refresh message supports the refreshing of RSVP state without the transmission of conventional Path messages. |
| Refresh RESV msg parsing | Refresh message supports the refreshing of RSVP state without the transmission of conventional Resv messages. |
| Detour identification | Detours are calculated to avoid the immediate downstream link and node. |

show rsvp admin-groups

Use this command to display all known administrative groups configured through the NSM for the system.

Command Syntax

```
show rsvp admin-groups
```

Parameters

None

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

This is a sample output showing four administrative groups configured through NSM.

```
#show rsvp admin-groups
Admin group detail:
  Value of 0 associated with admin group 'a'
  Value of 1 associated with admin group 'b'
  Value of 2 associated with admin group 'c'
  Value of 3 associated with admin group 'd'
#
```

[Table 36](#) explains the show command output fields.

Table 36. show rsvp admin-groups output field

| Field | Description |
|--------------------|---|
| Admin group detail | Administrative groups details which implements the link coloring of resource classes. |

show rsvp bypass

Use this command to display bypass session related information for configured bypass LSPs.

Command Syntax

```
show rsvp bypass
```

Parameters

None

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp bypass
Ingress RSVP:
To          From          Tun-ID  LSP-ID  LSPName          State
Uptime    Rt   Style  Labelin Labelout
172.31.54.4 172.31.54.1 5001    2201    BYPASS2-172.31.222.19-Bypass  UP    02d15h11m 1 1
SE        -    52516
172.31.54.2 172.31.54.1 5002    2202    BYPASS3-172.31.222.9-Bypass   UP    02d15h11m 1 1
SE        -    0
172.31.54.2 172.31.54.1 5003    2203    BYPASS4-172.31.222.7-Bypass   UP    02d15h11m 1 1
SE        -    0
172.31.53.18 172.31.54.1 5004    2204    BYPASS5-172.31.189.179-Bypass UP    02d15h11m 1 1
SE        -    52501
```

show rsvp bypass detail

Use this command to display bypass session related information in detail for all configured bypass LSPs or the bypass session with specified bypass tunnel name.

Command Syntax

```
show rsvp bypass (BYPASSNAME | detail)
```

Parameters

BYPASSNAME

Bypass tunnel name

detail

Detailed information of all configured bypass sessions

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp bypass BYPASS2-172.31.222.19
Ingress (Bypass)
172.31.54.4
  From: 172.31.54.1, LSPstate: Up, LSPname: BYPASS2-172.31.222.19-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 324ms
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
  IGP-Shortcut: Disabled, LSP metric: 1
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out: 52516,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 172.31.54.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 172.31.222.25, po22
  Path refresh: 30 seconds (RR enabled) (due in 12409 seconds)
  Resv lifetime: 157 seconds (due in 130 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: none
  Exclude Link: 172.31.222.19
  Session Explicit Route Detail :
    172.31.222.25/32 strict
    172.31.180.3/32 strict
    172.31.180.4/32 strict
  Record route:
```

```
-----  
IP Address      Label  
-----  
<self>  
172.31.222.25  
172.31.180.3  
172.31.180.4  
Style: Shared Explicit Filter  
Traffic type: controlled-load  
Minimum Path MTU: 9174  
Current Error:  
  Code : None, Value : None  
  Originated Node : None, Recorded Time : N/A  
Last Signaled Error:  
  Code : None, Value : None  
  Originated Node : None, Recorded Time : N/A  
Trunk Type: mpls  
Total LSP protected : 0, Bandwidth in use : 0
```

show rsvp bypass lsp-address-list

Use this command to display address details of every node of a bypass session shown as merge node detail for egress node of bypass session and transit node detail for transit node details of bypass session.

Command Syntax

```
show rsvp bypass (BYPASSNAME|) lsp-address-list
```

Parameters

BYPASSNAME

(Optional) Bypass tunnel name

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp bypass BYPASS2-172.31.222.19 lsp-address-list
Bypass trunk: BYPASS2-172.31.222.19

Merge Point Router ID: 172.31.54.4

Number of Merge Point IP addresses: 6
IP address:
  172.31.222.22    172.31.180.4    172.31.222.19    172.31.222.27
  172.31.222.31    172.31.186.4

Number of Transit Point IP addresses: 9
IP address:
  172.31.54.3      172.31.222.23    172.31.222.30    172.31.180.2
  172.31.222.25    172.31.186.20    172.31.33.120    172.31.180.3
  172.31.180.5

LSP address query interval: 60 seconds, next retry in: 27 seconds
```

show rsvp bypass protected-lsp-list

Use this command to display the list of sessions protected by a bypass session and match code provides the details bypass is a perfect match or any constraint compromised.

**Note:**

Match code 0 is an indication of perfect match i.e. all constraint of protected session matched. i.e. If protected session asked for node protection, then bypass provides perfect node protection by merging exactly at next to next hop node. If protected session asked for bandwidth protection, bypass provides bandwidth protection. In case of PHP node, even when node protection is requested by protected session, it is not applicable and node protection request is not applicable on PHP node. Thus, a bypass providing link protection with other criteria matching is considered as perfect match.

If a bypass protected session requested for link protection but it is mapped to a bypass node protection, then it is not a perfect match. Match code will be 4 in that case.

When bandwidth protection is requested, highest importance of bypass mapping given to bandwidth protection. When bandwidth protection cannot be provided, then the remaining constraints given importance.

Command Syntax

```
show rsvp bypass (BYPASSNAME|) protected-lsp-list
```

Parameters

BYPASSNAME

Bypass tunnel name

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp bypass protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
           3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: BYPASS2-172.31.222.19
Bypass trunk bandwidth type: best-effort
Total LSP protected : 0
Bandwidth in use : 0

Bypass trunk: BYPASS3-172.31.222.9
Bypass trunk bandwidth type: best-effort
```

List of LSP's Protected:

| Tunnel-id id | Lsp-Id Ingress | Lsp-Name Egress | Match-Code | Role | Ext_tnl_ | | |
|-------------------------|-------------------|------------------------------------|---------------|-------------|---------------|---------------|---|
| 61976 | 3 | to_OKL_ | | | | | |
| STRICT | | Transit | 172.31.2.52 | 172.31.2.52 | 172.31.54.2 | 0 | |
| 61975 | 4 | to_OKL_2ND_ | | | | | |
| LOOSE | | Transit | 172.31.2.52 | 172.31.2.52 | 172.31.54.2 | 0 | |
| 20 | 23884 | to_OKL_1ST_LOOSE::to_OKL_1ST_LOOSE | | | | | |
| | | Transit | | | 172.31.33.120 | 172.31.33.120 | 1 |
| 72.31.54.2 | 0 | | | | | | |
| 22 | 5478 | to_OKL_2ND_LOOSE::to_OKL_2ND_LOOSE | | | | | |
| | | Transit | | | 172.31.33.120 | 172.31.33.120 | 1 |
| 72.31.54.2 | 0 | | | | | | |
| 61974 | 3 | to_OKL_1ST_ | | | | | |
| LOOSE | | Transit | 172.31.2.52 | 172.31.2.52 | 172.31.54.2 | 0 | |
| 21 | 36172 | to_OKL_STRICT::to_OKL_ | | | | | |
| STRICT | Transit | 172.31.33.120 | 172.31.33.120 | 172.31.54.2 | 0 | | |
| Total LSP protected : 6 | | | | | | | |
| Bandwidth in use : 0 | | | | | | | |

show rsvp control-adjacency

Use this command to display RSVP specific information for control adjacency.

Command Syntax

```
show rsvp control-adjacency
show rsvp control-adjacency CANAME
```

Parameters

CANAME

Use this parameter to display the name of a control-adjacency

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#"show rsvp control-adjacency" without parameters:
Control Adj      Admin status  Oper Status    Peer-address   Gifindex       Control Channel

#"show rsvp control-adjacency" with parameters:
Admin Status    "Enabled" : "Disabled"
Oper Status     "Up" : "Down"
Peer-address
Gifindex
Control-Channel in use  cc->name : "N/A"
Control-Channel Gifindex
Control-Channel Local-address
Control-Channel Peer-address
Control-Channel ID
Control-Channel Binding Ifindex
Refresh Reduction usage  "Disabled" : "Enabled"
Message Acknowledgement "Enabled" : "Disabled"
Bundle Buffer size
Current Epoch Value
Primary IPv4 address  IPv4_address : "N/A"
Primary IPv6 address  IPv6_address : "N/A"
Configured refresh time
Configured keep multiplier
Acknowledgement Await Timeout
Hello Receipt      "Enabled" : "Disabled"
Hello Interval
Hello Timeout
Non IANA Hello exchange  "Enabled" : "Disabled"
```

[Table 37](#) explains the show command output fields.

Table 37. show rsvp control-adjacency output field

| Field | Description |
|------------------------------|---|
| Control Adj | Control Adjacency status and configuration. |
| Admin status | Indicates whether the user can administratively disable a peer while still preserving its configuration. Up = Yes, Down = No. |
| Oper Status | Displays the current status of the cross-connect segment – Up or Down. |
| Peer-address | Peer address in aa IPv4 and IPv6 format. |
| Gifindex | Number of gif index on which RSVP is active. |
| Control Channel | Control Channel status and configuration. |
| Refresh Reduction usage | Measure of processing over head requests of refresh messages. |
| Message Acknowledgment | The router that initiates the acknowledgment messages for an RSVP session. |
| Bundle Buffer size | Number of bundle buffer size. |
| Current Epoch Value | Value of the database epoch and number of entries in the epoch. |
| Primary IPv4 address | Primary IPv4 address of the neighbor interface. |
| Primary IPv6 address | Primary IPv6 address of the neighbor interface. |
| Configured Refresh Time | Time refresher which takes to generate periodic RSVP messages. |
| Configured Keep Multiplier | Number of RSVP messages that can be lost before an RSVP state is declared stale. |
| Acknowledgment Await Timeout | The router that initiates the acknowledgment messages for an RSVP session waits for the timeout. |
| Hello Receipt | To exchange Hello messages among neighbors. |
| Hello Interval | Frequency at which RSVP hellos are sent on this interface (in seconds). |
| Hello Timeout | RSVP Hello State Timer feature detects when a neighbor is down and triggers faster state timeout. |
| Non IANA Hello exchange | Hello exchange state in the interface. |

show rsvp data-link

Use this command to display RSVP specific information for data links.

Command Syntax

```
show rsvp data-link  
show rsvp data-link DLNAME
```

Parameters

DLNAME

Data link name

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#sh rsvp data-link
```

show rsvp graceful-restart

To modify the lines displayed, use the | (output modifier token); to save the output to a file, use the > output redirection token.

Command Syntax

```
show rsvp graceful-restart
show rsvp graceful-restart A.B.C.D
```

Parameters

A.B.C.D

IPv4 address of a specific neighbor (optional).

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 5.0.

Example

```
#show rsvp graceful-restart
Graceful Restart: Enabled
Advertised Restart Time: 180000 msec
Advertised Recovery Time: 360000 msec
Sending Recovery Time: Yes
Remote addr: 172.16.10.2 Local addr: 172.16.10.1
Nbr State: Normal Type: Reroute
Nbr Hello State: Up
LSPs protecting: 0
Restart Time: 0 msec, Recovery Time: 0 msec
Rest of Restart Time: 0 msec, Rest of Recovery Time: 0 msec
```

show rsvp interface

Use this command to display data about RSVP-specific information for interfaces, or about a specific interface.

Command Syntax

```
show rsvp interface
show rsvp interface IFNAME
```

Parameter

IFNAME

The name of the interface to display data.

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp interface eth0
Status                               : Enabled
Interface Index                      : 2
Refresh Reduction usage              : Enabled
Message Acknowledgement             : Disabled
Bundle Buffer size                   : 65535
Current Epoch Value                  : 208043005
Primary IPv4 address                 : 10.10.23.1
Primary IPv6 address                 : N/A
Interface Type                      : Ethernet
Administrative Group                 : a
                                     : d
Configured refresh time              : 5
Configured keep multiplier           : 3
Acknowledgement Await Timeout       : 10
Hello Receipt                       : Disabled
Hello Interval                      : 2
Hello Timeout                       : 10
Non IANA Hello exchange              : Disabled
#
```

[Table 38](#) explains the show command output fields.

Table 38. show rsvp interface output field

| Field | Description |
|-------------------------|---|
| Status | Display the status of Resource Reservation Protocol (RSVP). |
| Interface Index | Number of interface index on which RSVP is active. |
| Refresh Reduction usage | Measure of processing over head requests of refresh messages. |

Table 38. show rsvp interface output field (continued)

| Field | Description |
|------------------------------|---|
| Message Acknowledgement | The router that initiates the acknowledgment messages for an RSVP session. |
| Bundle Buffer size | Number of bundle buffer size. |
| Current Epoch Value | Value of the database epoch and number of entries in the epoch. |
| Primary IPv4 address | Primary IPv4 address of the neighbor interface. |
| Primary IPv6 address | Primary IPv6 address of the neighbor interface. |
| Interface Type | Type of interface. |
| Administrative Group | The administrators who belong to the same administrative group. |
| Configured Refresh Time | Time refresher which takes to generate periodic RSVP messages. |
| Configured Keep Multiplier | Number of RSVP messages that can be lost before an RSVP state is declared stale. |
| Acknowledgment Await Timeout | The router that initiates the acknowledgment messages for an RSVP session waits for the timeout. |
| Hello Receipt | To exchange Hello messages among neighbors. |
| Hello Interval | Frequency at which RSVP hellos are sent on this interface (in seconds). |
| Hello Timeout | RSVP Hello State Timer feature detects when a neighbor is down and triggers faster state timeout. |
| Non IANA Hello exchange | Hello exchange state in the interface. |

show rsvp l2-info

Use this command to display MAC and out interface details of a bypass tunnel which is used to send control messages of protected sessions over bypass tunnel when protected session is using backup.

Command Syntax

```
show rsvp l2-info
```

Parameters

None

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp l2-info
=====
## Bypass ftn l2 info ##
Ftn IX: 1
Out label: 52521 Out if 100022
src addr:(34ef.b63d.57a9)
Dst addr:(34ef.b694.3e08)
=====
## Bypass ftn l2 info ##
Ftn IX: 2
Out label: 3 Out if 100022
src addr:(34ef.b63d.57a9)
Dst addr:(34ef.b694.3e08)
=====
```

show rsvp local-addresses

Use this command to display data about any configured RSVP local address, including either IPv4 or IPv6 addresses.

Command Syntax

```
show rsvp local-addresses
show rsvp local-addresses ipv4
show rsvp local-addresses ipv6
```

Parameters

ipv4

Use this parameter to display IPv4 local addresses.

ipv6

Use this parameter to display IPv6 local addresses.

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp local-addresses
IPv4 Addresses:
Address                Interface
4.4.4.40               lo
10.1.1.2.40            eth0
14.14.14.8             eth4
34.0.0.40              eth2
80.0.0.40              eth2
127.0.0.1              lo
IPv6 Addresses:
Address                Interface
::1                   lo
fe80::202:b3ff:fed5:8dbb eth4
fe80::202:b3ff:fed5:9842 eth2
fe80::20e:cff:fe83:3727  eth0
#
```

[Table 39](#) explains the show command output fields.

Table 39. show rsvp local-addresses output field

| Field | Description |
|----------------|---------------------------------|
| IPv4 Addresses | IPv4 address for the interface. |
| IPv6 Addresses | IPv6 address for the interface. |
| Address | Address for the interface. |
| Interface | Name of the interface. |

show rsvp lsp-id-table

Use this command to display summary information for the RSVP LSP-ID allocation table.

The summary includes each block's index range, the number of usable and free indices, and the first available LSP-ID index. Total used and free indices across all allocated blocks are also shown.

Command Syntax

```
show rsvp lsp-id-table
show rsvp lsp-id-table block-id <BLOCK-ID>
```

Parameters

BLOCK-ID

Specifies the block-id to display detailed bitmap information for a specific LSP-ID block, including per-bit allocation status.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

Following is a sample output from the command using the detail parameter.

```
vpc2#show rsvp lsp-id-table
Bitmap Table: lsp-id, Range: 2201-65535, Current block: 1
+-----+-----+-----+-----+-----+
block_id  index_min  index_max  usable_indices  free_indices  first_free_index
+-----+-----+-----+-----+-----+
1          2201      2559        359           358           2202

Total - blocks: 1,  used-indices: 1,  free-indices: 358

vpc2#
```

Following is a sample output from the command using the BLOCK-ID parameter.

```
vpc2#show rsvp lsp-id-table block-id 1
Block ID           : 1
Bitmap Table       : lsp-id
Range              : 2201-2559
Usable indices     : 359
Free index count   : 358
First free index   : 2202
Flags              : NONE
Bitmap array (1 - used, 0 - free) :
[ 2176 ]                                1 0 0 0 0 0 0 [ 2207 ]
[ 2208 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2239 ]
[ 2240 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2271 ]
[ 2272 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2303 ]
[ 2304 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2335 ]
[ 2336 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2367 ]
```



```
[ 2368 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2399 ]  
[ 2400 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2431 ]  
[ 2432 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2463 ]  
[ 2464 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2495 ]  
[ 2496 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2527 ]  
[ 2528 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 2559 ]
```

vpc2#

show rsvp neighbor

Use this command to display a list of IPv4 RSVP neighbors or just a single IPv4 RSVP neighbor.

Command Syntax

```
show rsvp neighbor
show rsvp neighbor A.B.C.D
show rsvp neighbor detail
```

Parameters

A.B.C.D

Use this parameter to display the IP address of the IPv4 RSVP neighbor.

detail

Detailed view.

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3. The *detail* parameter is updated in OcNOS version 6.6.1

Example

```
OcNOS#show rsvp neighbor detail
50.0.0.2 :
  Upstream LSPs: 0, Downstream LSPs: 1
  Source Instance: 1284681321, Destination Instance: 1610048250, Interface: ge7
  Neighbor supports Refresh Reduction, next SRefresh transmission in: 6s
  Neighbor supports Hello message, next Hello transmission in: 2s, neighbor expiry in: 7s
  Create Time: 00:01:55, Up Time: 00:01:26

  Tunnel ID  LSP ID    Ingress          Egress           Type             Associate_
Time  Name
  5001      2201        5.5.5.5          32.32.32.32      Downstream       00:00:18
    t7-Primary
60.0.0.2 :
  Upstream LSPs: 0, Downstream LSPs: 1
  Source Instance: 376354693, Destination Instance: 1405769949, Interface: xe10
  Neighbor supports Refresh Reduction, next SRefresh transmission in: 14s
  Neighbor supports Hello message, next Hello transmission in: 3s, neighbor expiry in: 8s
  Create Time: 00:01:55, Up Time: 00:01:25

  Tunnel ID  LSP ID    Ingress          Egress           Type             Associate_
Time  Name
  5002      2202        5.5.5.5          32.32.32.32      Downstream       00:00:11
    B1-Bypass
OcNOS#
```

```
OcNOS#show rsvp neighbor 10.0.0.2
  Upstream LSPs: 0, Downstream LSPs: 1
  Source Instance: 1443372054, Destination Instance: 2075943448, Interface: xe2
  Neighbor supports Refresh Reduction, next SRefresh transmission in: 16s
```

```
Neighbor supports Hello message, next Hello transmission in: 3s, neighbor expiry in: 5s
Create Time: 00:02:19, Up Time: 00:00:46
```

| Tunnel ID | LSP ID | Ingress | Egress | Type | Associate_ |
|-----------|------------|---------|---------|------------|------------|
| Time Name | | | | | |
| 5001 | 2201 | 1.1.1.1 | 3.3.3.3 | Downstream | 00:02:19 |
| | t2-Primary | | | | |

OcNOS#

```
#show rsvp neighbor
IP Address      UpStrm LSP DnStrm LSP RefreshReduc Srefresh In  Type
10.10.20.4      0          1      Enabled      5s          Implicit
10.10.23.2      0          1      Enabled      8s          Implicit
#
```

[Table 40](#) explains the show command output fields.

Table 40. show rsvp neighbor output field

| Field | Description |
|---------------|---|
| IP Address | Address for the interface. |
| UpStrm LSP | Specify the upstream label for the bidirectional label-switched path (LSP). |
| DnStrm LSP | Specify the dnstream label for the bidirectional label-switched path (LSP). |
| Refresh Reduc | Refresh reduction improves the scalability, latency, and reliability of Resource Reservation Protocol (RSVP) signaling to enhance network performance and message delivery. |
| Srefresh In | Remaining seconds for srefresh timer expiry. |
| Type | Type of neighbor interface. |

show rsvp nexthop-cache

Use this command to display the current nexthops being cached by RSVP.

Command Syntax

```
show rsvp nexthop-cache
```

Parameters

None

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp nexthop-cache
Prefix      Nexthop      Outgoing Intf  Valid For    Num Sessions
10.10.20.80/32  0.0.0.0      eth1           12 seconds   1
10.10.23.60/32  0.0.0.0      eth0           17 seconds   1
#
```

[Table 41](#) explains the show command output fields.

Table 41. show rsvp nexthop-cache output field

| Field | Description |
|---------------|---|
| Prefix | It is an ordered list and entries are evaluated in order of increasing sequence number. |
| Nexthop | IP address of the next hop. |
| Outgoing Intf | Short name of the physical interface through which traffic goes to the protected link. |
| Valid For | Frequency at which RSVP hellos are sent next hop on this interface (in seconds). |
| Num Sessions | Number of session in the interface. |

show rsvp path

Use this command to display the configured rsvp paths and their configured hops. Specify the pathname to show hops related to a specific path. If no pathname is specified all the rsvp paths are displayed.

Command Syntax

```
show rsvp path
show rsvp path PATHNAME
```

Parameter

PATHNAME

The name of a specific path.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

Following are sample outputs from this command, with and without a **PATHNAME** (PRI) specified.

```
#show rsvp path
Path name: PRI, id: 1
 10.10.11.51 strict
 10.10.12.50 strict
 10.10.13.51 strict

Path name: SEC, id: 2
 10.10.10.51 strict

Path name: loop, id: 3
 10.10.11.51 strict
 10.10.12.50 strict
 10.10.13.51 strict
 10.10.14.50 strict
#

#show rsvp path PRI
Path name: PRI, id: 1
 10.10.11.51 strict
 10.10.12.50 strict
 10.10.13.51 strict
#
```

[Table 42](#) explains the show command output fields.

Table 42. show rsvp path output field

| Field | Description |
|-----------|---------------------------|
| Path name | Name of the path. |
| id | Address of the rsvp path. |

show rsvp protected-lsp-reop-list

Use this command to display list of facility protected sessions which didn't get any bypass protection or didn't get a perfect bypass protection. These sessions are checked for better protection whenever a new bypass session comes up.

Command Syntax

```
show rsvp protected-lst-reop-list
```

Parameters

None

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp protected-lsp-reop-list
Tunnel-id  Lsp-Id    Lsp-Name      Role      Ext_tnl_
id          Ingress      Egress        Protected
222         169          LHR_          Transit   172.31.53.18    172.31.53.18    172.31.2.52     Yes
t222
204         1522         LHR_          Transit   172.31.53.18    172.31.53.18    172.31.33.120   Yes
t204
17          52608        GGN_NDLS_2ND_LOOSE::to_CISCO_2ND_LOOSE
72.31.53.18  Yes          Transit       172.31.33.120   172.31.33.120   1
```

show rsvp session

Use this command to display session-related information for configured LSPs.

Command Syntax

```
show rsvp session
show rsvp session up
show rsvp session up detail
show rsvp session down
show rsvp session down detail
```

Parameters

up

Use this parameter to display sessions that are currently operational.

down

Use this parameter to display sessions that are currently not operational.

detail

Use this parameter to display detailed session-related information.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

Following is a sample output from the command using the detail parameter.

```
#show rsvp session detail
Ingress (Primary)
10.10.21.3
  From: 1.1.1.1, LSPstate: Up, LSPname: t1
  Setup priority: 5, Hold priority: 5
  CSPF usage: Disabled
  LDP Tunneling : Disabled
  LSP Protection: None
  Label in: -, Label out: 16,
  Tspecc rate: 10m, Fspec rate: 10m
  Tunnel Id: 1, LSP Id: 2, Ext-Tunnel Id: 1.1.1.1
  Downstream: 10.10.23.2, eth0
  Path refresh: 5 seconds (due in 6772 seconds)
  Resv lifetime: 26 seconds (due in 25 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: p1 (in use)
  Configured Explicit Route Detail :
    10.10.23.2/32 strict
  Session Explicit Route Detail :
    10.10.23.2/32 strict
  Record route: <self> 10.10.23.2 10.10.21.3
```

```

Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_SIGNAL
CLASS      DSCP_value      EXP_value
#

```

Here is the explanation of the "show command" output fields.

Table 43. "show rsvp session detail "output field

| Field | Description |
|-------------------|---|
| Ingress (Primary) | Information about ingress RSVP sessions. Each session has one line of output. |
| From | Source (ingress switch) of the session. |
| LSP state | State of the LSP that is being handled by this RSVP session. It can be either Up, Dn (down), or Admin Dn. Admin Dn indicates that the LSP is being taken down gracefully. |
| LSPname | Name of the LSP. |
| Setup priority | Value of the setup priority. |
| Hold priority | Determines the degree to which an LSP holds onto its session reservation after the LSP has been set up successfully. |
| CSPF usage | CSPF usage state in the rsvp session. |
| LDP Tunnelling | Displays the status of LDP tunneling. In the given RSVP session, the status can be either enabled or disabled for LDP over RSVP. |
| LSP Protection | Protects the traffic failures. |
| Label in | Incoming label for this LSP. |
| Label out | Outgoing label for this LSP. |
| Tspec rate | Sender's traffic specification, which describes the sender's traffic parameters. |
| Fspec rate | Fspec peak rate values. |
| Tunnel id | Tunnel address (destination port) for the session. |
| LSP id | Address of the LSP in the interface. |
| Ext-Tunnel Id | Session address for the ext-tunnel. |
| Down stream | Specify the dstream label for the bidirectional label-switched path (LSP). |
| Path refresh | Path messages are sent periodically to refresh path states. The refresh interval is controlled by a variable called the refresh time. |
| Resv lifetime | Number of seconds remaining in the lifetime of the reservation. |
| Retry count | Number of times sanity polling periodically checks for an error condition in the FPC. |
| | Interval sets the time for the messages in order to control the session. |
| LSP Type | Type of ELSP signal. |

```

#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```


Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|------------------------------------|----------|----------|---------|----------|----------------------------------|-----------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.1 | 30.0.1.2 | 5001 | 2201 | PRI | PE3-PE- | |
| Primary | UP | 01:13:41 | 1 1 SE | - | 24332 | |
| 30.0.1.4 | 30.0.1.2 | 5002 | 2202 | PRI | PE3-PE2- | |
| Primary | UP | 01:14:11 | 1 1 SE | - | 24328 | |
| 30.0.1.7 | 30.0.1.2 | 5003 | 2203 | PRI | PE3-PE4- | |
| Primary | UP | 01:05:41 | 1 1 SE | - | 24324 | |
| 30.0.1.5 | 30.0.1.2 | 5004 | 2204 | PRI | PE3-RR2- | |
| Primary | UP | 01:14:09 | 1 1 SE | - | 24321 | |
| 30.0.1.5 | 30.0.1.2 | 5004 | 2209 | SEC | PE3-RR2- | |
| Secondary | UP | 00:00:04 | 1 1 SE | - | 24325 | |
| 30.0.1.1 | 30.0.1.2 | 5005 | 2205 | BPS | nhop- | |
| Bypass | UP | 01:06:05 | 1 1 SE | - | 24321 | |
| 30.0.1.4 | 30.0.1.2 | 5006 | 2206 | BPS | nnhop- | |
| Bypass | DN | N/A | 0 0 SE | - | - | |
| 30.0.1.4 | 30.0.1.2 | 5007 | 2207 | BPS | BL-10.0.1.19-30.0.1.4-100-Bypass | UP 01:14: |
| 10 1 1 SE | - | 24327 | | | | |
| 30.0.1.4 | 30.0.1.2 | 5008 | 2208 | BPS | BN-30.0.1.1-30.0.1.4-101-Bypass | |
| UP 01:14:09 | 1 1 SE | - | 3 | | | |
| 30.0.1.5 | 30.0.1.2 | 5010 | 2210 | BPS | BL-10.0.1.23-30.0.1.5-102-Bypass | UP 01:13: |
| 41 1 1 SE | - | 24325 | | | | |
| 30.0.1.3 | 30.0.1.2 | 5011 | 2211 | BPS | BL-10.0.1.5-30.0.1.3-103-Bypass | |
| UP 01:05:41 | 1 1 SE | - | 24330 | | | |
| Total 11 displayed, Up 10, Down 1. | | | | | | |

Transit RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|----------------------------------|----------|----------|---------|----------|----------------------------------|-----------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.1 | 30.0.1.7 | 5001 | 2201 | PRI | PE4-PE1- | |
| Primary | UP | 01:13:40 | 1 1 SE | 24324 | 24324 | |
| 30.0.1.7 | 30.0.1.4 | 5003 | 2210 | PRI | PE2-PE4- | |
| Secondary | UP | 01:10:43 | 1 1 SE | 24326 | 24328 | |
| 30.0.1.4 | 30.0.1.7 | 5006 | 2209 | PRI | BL-10.0.1.32-30.0.1.4-100-Bypass | UP 01:12: |
| 38 1 1 SE | 24325 | 3 | | | | |
| 30.0.1.5 | 30.0.1.4 | 5006 | 2207 | PRI | BL-10.0.1.13-30.0.1.5-100-Bypass | UP 01:14: |
| 10 1 1 SE | 24320 | 3 | | | | |
| 30.0.1.7 | 30.0.1.1 | 5006 | 2206 | PRI | PE1-PE4_3- | |
| Primary | UP | 01:05:47 | 1 1 SE | 24331 | 24323 | |
| Total 5 displayed, Up 5, Down 0. | | | | | | |

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-----------|----------|---------|----------|----------------------------------|-----------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.2 | 30.0.1.1 | 5001 | 2201 | PRI | PE1-PE3- | |
| Primary | UP | 01:14:05 | 1 1 SE | 24321 | - | |
| 30.0.1.2 | 30.0.1.1 | 5001 | 2210 | PRI | PE1-PE3- | |
| Secondary | UP | 01:14:05 | 1 1 SE | 24322 | - | |
| 30.0.1.2 | 30.0.1.5 | 5003 | 2209 | PRI | RR2-PE3- | |
| Secondary | UP | 00:01:29 | 1 1 SE | 24323 | - | |
| 30.0.1.2 | 30.0.1.5 | 5003 | 2210 | PRI | RR2-PE3- | |
| Primary | UP | 00:01:29 | 1 1 SE | 24332 | - | |
| 30.0.1.2 | 30.0.1.7 | 5003 | 2203 | PRI | PE4-PE3- | |
| Primary | UP | 01:06:07 | 1 1 SE | 24328 | - | |
| 30.0.1.2 | 10.0.1.29 | 5003 | 2203 | PRI | PE4-PE3- | |
| Detour | UP | 01:06:07 | 1 1 SE | 24330 | - | |
| 30.0.1.2 | 10.0.1.30 | 5003 | 2203 | PRI | PE4-PE3- | |
| Detour | UP | 01:06:07 | 1 1 SE | 24329 | - | |
| 30.0.1.2 | 30.0.1.4 | 5007 | 2208 | PRI | BN-30.0.1.1-30.0.1.2-101-Bypass | |
| UP 01:13:59 | 1 1 SE | 3 | - | | | |
| 30.0.1.2 | 30.0.1.5 | 5007 | 2208 | PRI | BL-10.0.1.22-30.0.1.2-101-Bypass | UP 00:01: |
| 37 1 1 SE | 3 | - | | | | |
| 30.0.1.2 | 30.0.1.1 | 5012 | 2217 | PRI | BL-10.0.1.3-30.0.1.2-101-Bypass | |
| UP 01:13:59 | 1 1 SE | 3 | - | | | |

Total 10 displayed, Up 10, Down 0.

PE3-RR1-7003#

PE3-RR1-7003#show rsvp session

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass

State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary

* indicates the session is active with local repair at one or more nodes

(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|----------|----------|---------|----------|----------------------------------|-----------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.1 | 30.0.1.2 | 5001 | 2201 | PRI | PE3-PE- | |
| Primary | UP | 01:14:15 | 1 1 SE | - | 24332 | |
| 30.0.1.4 | 30.0.1.2 | 5002 | 2202 | PRI | PE3-PE2- | |
| Primary | UP | 01:14:45 | 1 1 SE | - | 24328 | |
| 30.0.1.7 | 30.0.1.2 | 5003 | 2203 | PRI | PE3-PE4- | |
| Primary | UP | 01:06:15 | 1 1 SE | - | 24324 | |
| 30.0.1.5 | 30.0.1.2 | 5004 | 2204 | PRI | PE3-RR2- | |
| Primary | UP | 01:14:43 | 1 1 SE | - | 24321 | |
| 30.0.1.5 | 30.0.1.2 | 5004 | 2209 | SEC | PE3-RR2- | |
| Secondary | UP | 00:00:38 | 1 1 SE | - | 24325 | |
| 30.0.1.1 | 30.0.1.2 | 5005 | 2205 | BPS | nhop- | |
| Bypass | UP | 01:06:39 | 1 1 SE | - | 24321 | |
| 30.0.1.4 | 30.0.1.2 | 5006 | 2206 | BPS | nnhop- | |
| Bypass | DN | N/A | 0 0 SE | - | - | |
| 30.0.1.4 | 30.0.1.2 | 5007 | 2207 | BPS | BL-10.0.1.19-30.0.1.4-100-Bypass | UP 01:14: |
| 44 1 1 SE | - | 24327 | | | | |
| 30.0.1.4 | 30.0.1.2 | 5008 | 2208 | BPS | BN-30.0.1.1-30.0.1.4-101-Bypass | |
| UP 01:14:43 | 1 1 SE | - | 3 | | | |
| 30.0.1.5 | 30.0.1.2 | 5010 | 2210 | BPS | BL-10.0.1.23-30.0.1.5-102-Bypass | UP 01:14: |
| 15 1 1 SE | - | 24325 | | | | |
| 30.0.1.3 | 30.0.1.2 | 5011 | 2211 | BPS | BL-10.0.1.5-30.0.1.3-103-Bypass | |
| UP 01:06:15 | 1 1 SE | - | 24330 | | | |

Total 11 displayed, Up 10, Down 1.

Transit RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-----------|----------|----------|---------|----------|----------------------------------|-----------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.1 | 30.0.1.7 | 5001 | 2201 | PRI | PE4-PE1- | |
| Primary | UP | 01:14:14 | 1 1 SE | 24324 | 24324 | |
| 30.0.1.7 | 30.0.1.4 | 5003 | 2210 | PRI | PE2-PE4- | |
| Secondary | UP | 01:11:17 | 1 1 SE | 24326 | 24328 | |
| 30.0.1.4 | 30.0.1.7 | 5006 | 2209 | PRI | BL-10.0.1.32-30.0.1.4-100-Bypass | UP 01:13: |
| 12 1 1 SE | 24325 | 3 | | | | |
| 30.0.1.5 | 30.0.1.4 | 5006 | 2207 | PRI | BL-10.0.1.13-30.0.1.5-100-Bypass | UP 01:14: |
| 44 1 1 SE | 24320 | 3 | | | | |
| 30.0.1.7 | 30.0.1.1 | 5006 | 2206 | PRI | PE1-PE4_3- | |
| Primary | UP | 01:06:21 | 1 1 SE | 24331 | 24323 | |

Total 5 displayed, Up 5, Down 0.

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-----------|-----------|----------|---------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 30.0.1.2 | 30.0.1.1 | 5001 | 2201 | PRI | PE1-PE3- | |
| Primary | UP | 01:14:39 | 1 1 SE | 24321 | - | |
| 30.0.1.2 | 30.0.1.1 | 5001 | 2210 | PRI | PE1-PE3- | |
| Secondary | UP | 01:14:39 | 1 1 SE | 24322 | - | |
| 30.0.1.2 | 30.0.1.5 | 5003 | 2209 | PRI | RR2-PE3- | |
| Secondary | UP | 00:02:03 | 1 1 SE | 24323 | - | |
| 30.0.1.2 | 30.0.1.5 | 5003 | 2210 | PRI | RR2-PE3- | |
| Primary | UP | 00:02:03 | 1 1 SE | 24332 | - | |
| 30.0.1.2 | 30.0.1.7 | 5003 | 2203 | PRI | PE4-PE3- | |
| Primary | UP | 01:06:41 | 1 1 SE | 24328 | - | |
| 30.0.1.2 | 10.0.1.29 | 5003 | 2203 | PRI | PE4-PE3- | |
| Detour | UP | 01:06:41 | 1 1 SE | 24330 | - | |

```

30.0.1.2      10.0.1.30      5003      2203      PRI      PE4-PE3-
Detour      UP      01:06:41      1 1 SE      24329      -
30.0.1.2      30.0.1.4      5007      2208      PRI      BN-30.0.1.1-30.0.1.2-101-Bypass
UP      01:14:33      1 1 SE      3      -
30.0.1.2      30.0.1.5      5007      2208      PRI      BL-10.0.1.22-30.0.1.2-101-Bypass
                                                    UP      00:02:
11  1 1 SE      3      -
30.0.1.2      30.0.1.1      5012      2217      PRI      BL-10.0.1.3-30.0.1.2-101-Bypass
UP      01:14:33      1 1 SE      3      -
Total 10 displayed, Up 10, Down 0.

```

Here is the explanation of the "show command" output fields.

Table 44. Description of the “show rsvp session” output fields

| Field | Description |
|----------|--|
| To | Displays the destination IP address of the RSVP session. |
| From | Displays the source IP address of the RSVP session. |
| Tun-ID | Identifies the RSVP tunnel ID. |
| LSP-ID | Identifies the Label Switched Path (LSP) ID under the tunnel. |
| Type | Shows the session type. This information is applicable only for the ingress sessions. Since type is not signaled information non-ingress sessions won't have the information of the type. <ul style="list-style-type: none"> • PRI = Primary • SEC = Secondary • DTR = Detour • BPS = Bypass |
| LSPName | Displays the name of the RSVP LSP or session. The session name is derived from the trunk name by appending suffixes such as “-Primary, -Secondary, -Detour”, etc. |
| State | Shows the operational state of the session: <ul style="list-style-type: none"> • UP = Session is up • DN = Session is down • BU = Backup in use • SU = Secondary in use • FS = Forced to Secondary • UP* = Session is UP due to local repair in one or more downstream PLRs. |
| Uptime | Displays how long the session has been active. Uptime considers active sessions even in UP*, BU, SU states. |
| Rt | Indicates the QoS reservation for the session and the count of shared QoS reservation resources. |
| Style | Shows the RSVP reservation style (for example, SE = Shared Explicit). |
| Labelin | Displays the incoming MPLS label assigned to the session. |
| Labelout | Displays the outgoing MPLS label assigned to the session. |

show rsvp session count

Use this command to display session-related information for configured LSPs.

Command Syntax

```
show rsvp session count
show rsvp session count egress
show rsvp session count ingress
show rsvp session count transit
```

Parameters

egress

Use this parameter to display the number of configured egress sessions.

ingress

Use this parameter to display the number of configured ingress sessions.

transit

Use this parameter to display the number of configured transmit sessions.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp session count
Total configured: 1520, Up 1520, Down 0
#
```

[Table 45](#) explains the show command output fields.

Table 45. show rsvp session count output field

| Field | Description |
|------------------|---|
| Total configured | Number of configured rsvp session in the interface. |

show rsvp session egress

Use this command to display session-related information for an egress router.

Command Syntax

```
show rsvp session egress
show rsvp session egress A.B.C.D
show rsvp session egress X:X::X:X
show rsvp session egress detail
show rsvp session egress down
show rsvp session egress down detail
show rsvp session egress up
show rsvp session egress up detail
```

Parameters

A.B.C.D

Use this parameter to display an IPv4 address of an egress router

X:X::X:X

Use this parameter to display an IPv6 address of an egress router

down

Use this parameter to display sessions that are currently not operational

up

Use this parameter to display sessions that are currently operational

detail

Use this parameter to display detailed session-related information

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show rsvp session egress without parameters or with "up" or "down":
%s RSVP:
To                From                State                Pri Rt  Style Labelin Labelout
LSPName                               Uptime  Est.time  DStype
...
Total %d displayed

#show rsvp session egress with parameters:
"Bypass", "Primary", "Detour", "Secondary"
Make-Before-Break Sibling for session with LSP-ID:  prefix4: prefix6
From: u.prefix4: u.prefix6
LSPstate: %s, LSPname:
    "Up/"Using Backup"/"Using Secondary"
    "Dn",
Revert hold timer is ON due to expire in %d seconds
Revert Timer Finished, Forced Switch to Secondary LSP In Effect
CSPF usage: "Disabled" : "Enabled"
, CSPF Retry Count: %d, CSPF Retry Interval: %d seconds"
```

```

IGP-Shortcut: Enabled, LSP metric:
IGP-Shortcut: Disabled, LSP metric:
LSP Protection:
Bypass trunk:
Label in:
Label out:
Tspec rate:
Fspec rate:
Policer: Configured
    and installed in hardware
    but not installed in hardware
    Not Configured
Tunnel Id: %d, LSP Id: %d
Ext-Tunnel Id:
Downstream:
Upstream:
Path refresh: %d seconds (RR enabled), (due in %d seconds)
Path lifetime: %d seconds (due in %d seconds)
Resv refresh: %d seconds (due in %d seconds)
Resv lifetime: %d seconds (due in %d seconds)
Retry count: %d, intrvl: %d seconds", # remaining, next retry in: %d seconds",
RRO re-use as ERO: "Enabled" : "Disabled"
Label Recording: "Enabled" : "Disabled"
FRR Admin Groups/Admin Groups:
    ***admin group info***
Exclude path detail:
Exclude "Link" : "Node
Configured Path: "none" : "in use" : "not in use"
%s Explicit Route Detail "Configured" : "Received"
    "strict" : "loose"
Record route: " <self>") " ...incomplete"
Style: %s\n", rsvp_style_to_str (style));
Traffic type: "guaranteed" : "controlled-load" : "none"
Minimum Path MTU:
Traffic type: N/A
Minimum Path MTU: N/A
LSP Type: "ELSP_SIGNAL" : "ELSP_CONFIG"
CLASS DSCP_value EXP_value
The class to exp bits mapping is invalid.
LSP Type: L-LSP
LLSP DSCP: %d%d%d%d%d CLASS: %4s",
DSTE Class Type Number: Invalid, Class Type name(configured):
DSTE Class Type Number: %d, Class Type name:
Last Recorded Error Code: %s (%d)
Last Recorded Error Value: %s (%d)
Node where Last Recorded Error originated:
Trunk Type: "gmpls" : "mpls"
Tsid:
Merge Point Addresss [%d] =

```

[Table 46](#) explains the show command output fields.

Table 46. show rsvp session egress output field

| Field | Description |
|---------------------|---|
| LSP state | State of the LSP that is being handled by this RSVP session. It can be either Up, Dn (down), or Admin Dn. Admin Dn indicates that the LSP is being taken down gracefully. |
| LSP name | Name of the LSP. |
| CSPF usage | CSPF usage state in the rsvp session. |
| CSPF Retry Count | Number of times CSPF tried to find the path. |
| CSPF Retry Interval | The interval at which CSPF retry to find the path. |

Table 46. show rsvp session egress output field (continued)

| Field | Description |
|-------------------------------|---|
| IGP-Shortcut | Status of IGP shortcut for the RSVP trunk. |
| LSP metric | Relative/Absolute metric value of the LSP. |
| LSP Protection | LSP Protection configured for the RSVP trunk. |
| Bypass trunk | Name for the configured Bypass trunk. |
| Tspec rate | Sender's traffic specification, which describes the sender's traffic parameters. |
| Fspec rate | Fspec peak rate values. |
| Policer | QoS Policy configured for the RSVP trunk. |
| Tunnel Id | Tunnel identifier (destination port) for the RSVP session. |
| LSP Id | Address of the LSP in the interface. |
| Ext-Tunnel Id | Ext Tunnel identifier (destination port) for the RSVP session. |
| Down stream | Specify the dn stream label for the bidirectional label-switched path (LSP). |
| Upstream | Address of the previous hop for the egress session. |
| Path refresh | Path messages are sent periodically to refresh path states. The refresh interval is controlled by a variable called the refresh time. |
| Path lifetime | Number of seconds remaining in the lifetime of the reservation. |
| Resv refresh | Remaining time in seconds for the next Resv refresh. |
| Resv lifetime | Number of seconds remaining in the lifetime of the reservation. |
| Retry count | Number of times sanity polling periodically checks for an error condition in the FPC. |
| intrvl | Interval sets the time for the messages in order to control the session. |
| next retry in | Remaining time in seconds for the next retry. |
| RRO re-use as ERO | Enabling to re-use Record route as Explicit route for rsvp session. |
| Label Recording | Enabling to record the labels exchanged by all the peers. |
| FRR Admin Groups/Admin Groups | Resource affinities associated with the rsvp session. |
| Exclude path detail | Detailed List of the link addresses to be excluded for RSVP Bypass session. |
| Exclude Link | Address of the Link to be excluded for RSVP Bypass session. |
| Configured Path | Configured path name associated with the rsvp session. |
| Record route | Established rsvp path with each hop information. |
| Style | Reservation style associated with the rsvp session. |
| Traffic type | Traffic type associated with the rsvp session. |
| Minimum Path MTU | Path maximum transmission unit (MTU) discovery in the interface. |
| LSP Type | Type of ELSP signal. |
| CLASS | Name of the class which is associated with rsvp session. |
| DSCP_value | DSCP value of diff-serv class which is associated with rsvp session. |
| EXP_value | EXP value of diff-serv class which is associated with rsvp sess |

Table 46. show rsvp session egress output field (continued)

| Field | Description |
|---|--|
| DSTE Class Type Number | Diff-serv class type number associated with rsvp session. |
| Class Type name | Diff-serv class type name associated with rsvp session. |
| Last Recorded Error Code | The last recorded error code for the RSVP session. |
| Last Recorded Error Value | The last recorded error for the RSVP session. |
| Node where Last Recorded Error originated | Error originated node in the rsvp session. |
| Trunk Type | Trunk type in the rsvp session. |
| Tesid | Traffic Engineering Service Instance Identifier |
| Merge Point Address | Address of the node where the Bypass LSP joins with the protected LSP. |

show rsvp session ingress

Use this command to display session-related information for an ingress router.

Command Syntax

```
show rsvp session ingress
show rsvp session ingress A.B.C.D
show rsvp session ingress X:X::X:X
show rsvp session ingress detail
show rsvp session ingress down
show rsvp session ingress down detail
show rsvp session ingress up
show rsvp session ingress up detail
```

Parameters

A.B.C.D

Use this parameter to display an IPv4 address of an ingress router

X:X::X:X

Use this parameter to display an IPv6 address of an ingress router.

down

Use this parameter to display sessions that are currently not operational

up

Use this parameter to display sessions that are currently operational

detail

Use this parameter to display detailed session-related information

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

#show rsvp session ingress (without parameters or with up or down:)

```
%s RSVP:
To          From          State          Pri Rt  Style Labelin Labelout
LSPName          Uptime  Est.time      DStype
...
Total %d displayed

#show rsvp session ingress detail
Ingress (Primary)
41.41.41.31
From: 29.29.29.29, LSPstate: Up, LSPname: t1-Primary
Ingress FSM state: Operational
Establishment Time: 0s 3ms
Setup priority: 7, Hold priority: 0
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
IGP-Shortcut: Disabled, LSP metric: 1
```

```
LSP Protection: None
Label in: -, Label out: 24320, ELC
```

#show rsvp session ingress (with parameters:)

```
"Bypass", "Primary", "Detour", "Secondary"
Make-Before-Break Sibling for session with LSP-ID:   prefix4: prefix6
  From: u.prefix4: u.prefix6
LSPstate: %s, LSPname:
  "Up/"Using Backup"/"Using Secondary"
  "Dn",
  Revert hold timer is ON due to expire in %d seconds
  Revert Timer Finished, Forced Switch to Secondary LSP In Effect
  CSPF usage: "Disabled" : "Enabled"
  CSPF Retry Count: %d, CSPF Retry Interval: %d seconds"
  IGP-Shortcut: Enabled, LSP metric:
  IGP-Shortcut: Disabled, LSP metric:
  LSP Protection:
  Bypass trunk:
  Label in:
  Label out:
  Tspec rate:
  Fspec rate:
  Policer: Configured
    and installed in hardware
    but not installed in hardware
    Not Configured
  Tunnel Id: %d, LSP Id: %d
Ext-Tunnel Id:
  Downstream:
  Upstream:
  Path refresh: %d seconds (RR enabled), (due in %d seconds)
  Path lifetime: %d seconds (due in %d seconds)
  Resv refresh: %d seconds (due in %d seconds)
  Resv lifetime: %d seconds (due in %d seconds)
  Retry count: %d, intrvl: %d seconds", # remaining, next retry in: %d seconds",
  RRO re-use as ERO: "Enabled" : "Disabled"
  Label Recording: "Enabled" : "Disabled"
  FRR Admin Groups/Admin Groups:
    ***admin group info***
  Exclude path detail:
  Exclude "Link" : "Node
  Configured Path: "none" : "in use" : "not in use"
  %s Explicit Route Detail "Configured" : "Received"
    "strict" : "loose"
  Record route: " <self>)" " ...incomplete"
  Style: %s\n", rsvp_style_to_str (style));
  Traffic type: "guaranteed" : "controlled-load" : "none"
  Minimum Path MTU:
  Traffic type: N/A
  Minimum Path MTU: N/A
  LSP Type: "ELSP_SIGNAL" : "ELSP_CONFIG"
  CLASS      DSCP_value      EXP_value
  The class to exp bits mapping is invalid.
  LSP Type: L-LSP
  LLSP DSCP: %d%d%d%d%d%d CLASS: %4s",
  DSTE Class Type Number: Invalid, Class Type name(configured):
  DSTE Class Type Number: %d, Class Type name:
  Last Recorded Error Code: %s (%d)
  Last Recorded Error Value: %s (%d)
  Node where Last Recorded Error originated:
  Trunk Type: "gmpls" : "mpls"
  Tesid:
  Merge Point Address [%d] =
```

```
#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, B - BGP FTN, K - CLI FTN, t - tunnel
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code      FEC          FTN-ID Tunnel-id    Pri  LSP-Type    Out- Label  ELC  Out-Intf  Nexthop
R(t)> 2  9.29.29.29/32  1      5001        Yes  LSP_DEFAULT  24322    yes
eth2      1.41.41.31
R(t)> 29.29.29.29/32  2      5001        No   LSP_DEFAULT  4322     yes
eth1      69.69.69.42
```

[Table 47](#) explains the show command output fields.

Table 47. show rsvp session ingress output field

| Field | Description |
|---------------------|---|
| LSP state | State of the LSP that is being handled by this RSVP session. It can be either Up, Dn (down), or Admin Dn. Admin Dn indicates that the LSP is being taken down gracefully. |
| LSP name | Name of the LSP. |
| CSPF usage | CSPF usage state in the rsvp session. |
| CSPF Retry Count | Number of times CSPF tried to find the path. |
| CSPF Retry Interval | The interval at which CSPF retry to find the path. |
| IGP-Shortcut | Status of IGP shortcut for the RSVP trunk. |
| LSP metric | Relative/Absolute metric value of the LSP. |
| LSP Protection | LSP Protection configured for the RSVP trunk. |
| Bypass trunk | Name for the configured Bypass trunk. |
| Tspec rate | Sender's traffic specification, which describes the sender's traffic parameters. |
| Fspec rate | Fspec peak rate values. |
| Policer | QoS Policy configured for the RSVP trunk. |
| Tunnel Id | Tunnel identifier (destination port) for the RSVP session. |
| LSP Id | Address of the LSP in the interface. |
| Ext-Tunnel Id | Ext Tunnel identifier (destination port) for the RSVP session. |
| Down stream | Specify the dn stream label for the bidirectional label-switched path (LSP). |
| Upstream | Address of the previous hop for the egress session. |
| Path refresh | Path messages are sent periodically to refresh path states. The refresh interval is controlled by a variable called the refresh time. |
| Path lifetime | Number of seconds remaining in the lifetime of the reservation. |
| Resv refresh | Remaining time in seconds for the next Resv refresh. |
| Resv lifetime | Number of seconds remaining in the lifetime of the reservation. |
| Retry count | Number of times sanity polling periodically checks for an error condition in the FPC. |
| intrvl | Interval sets the time for the messages in order to control the session. |
| next retry in | Remaining time in seconds for the next retry. |
| RRO re-use as ERO | Enabling to re-use Record route as Explicit route for rsvp session. |

Table 47. show rsvp session ingress output field (continued)

| Field | Description |
|---|---|
| Label Recording | Enabling to record the labels exchanged by all the peers. |
| FRR Admin Groups/Admin Groups | Resource affinities associated with the rsvp session. |
| Exclude path detail | Detailed List of the link addresses to be excluded for RSVP Bypass session. |
| Exclude Link | Address of the Link to be excluded for RSVP Bypass session. |
| Configured Path | Configured path name associated with the rsvp session. |
| Record route | Established rsvp path with each hop information. |
| Style | Reservation style associated with the rsvp session. |
| Traffic type | Traffic type associated with the rsvp session. |
| Minimum Path MTU | Path maximum transmission unit (MTU) discovery in the interface. |
| LSP Type | Type of ELSP signal. |
| CLASS | Name of the class which is associated with rsvp session. |
| DSCP_value | DSCP value of diff-serv class which is associated with rsvp session. |
| EXP_value | EXP value of diff-serv class which is associated with rsvp sess |
| DSTE Class Type Number | Diff-serv class type number associated with rsvp session. |
| Class Type name | Diff-serv class type name associated with rsvp session. |
| Last Recorded Error Code | The last recorded error code for the RSVP session. |
| Last Recorded Error Value | The last recorded error for the RSVP session. |
| Node where Last Recorded Error originated | Error originated node in the rsvp session. |
| Trunk Type | Trunk type in the rsvp session. |
| Tesid | Traffic Engineering Service Instance Identifier. |
| Merge Point Addresss | Address of the node where the Bypass LSP joins with the protected LSP. |

show rsvp session LSP-NAME

Use this command to display information only for sessions with a specified name.



Note: This command doesn't work for sessions with tunnel name larger than 32 characters or sessions originated from non-OcNOS solutions.

Command Syntax

```
show rsvp session LSP-NAME
show rsvp session LSP-NAME primary
show rsvp session LSP-NAME secondary
```

Parameters

primary

Use this parameter to display any primary LSP sessions

secondary

Use this parameter to display any secondary LSP sessions

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

Following is a sample output from the command displaying session information about the LSP named t1.

```
#show rsvp session t1
Ingress (Primary)
192.168.0.90
  From: 192.168.0.63, LSPstate: Up, LSPname: t1
  Setup priority: 7, Hold priority: 0
  CSPF usage: Disabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  Label in: -, Label out: 17,
  Tspec rate: 0
  Tunnel Id: 1, LSP Id: 1, Ext-Tunnel Id: 192.168.0.63
  Downstream: 10.10.23.60, eth0
  Path refresh: 30 seconds (due in 34 seconds)
  Resv lifetime 157 seconds (due in 155 seconds)
  Retry Count: 0, Retry Interval: 30 seconds
  RRO re-use as ERO: Enabled
  Labels Recording: Disabled
  Admin Groups: include-any --> 0(a)
  Configured Path: p1 (in use)
  Configured Explicit Route Detail :
    10.10.23.60/32 loose
  Session Explicit Route Detail :
    10.10.23.60/32 loose
    10.10.21.90/32 loose
  Record route: <self> 10.10.23.60 10.10.21.90
```

```

Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Last Recorded Error Code: None
Last Recorded Error Value: None
#

```

[Table 48](#) explains the show command output fields.

Table 48. show rsvp session LSP-NAME output field

| Field | Description |
|----------------------------------|---|
| Ingress (Primary) | Information about ingress RSVP sessions. Each session has one line of output. |
| From | Source (ingress switch) of the session. |
| LSP state | State of the LSP that is being handled by this RSVP session. It can be either Up, Dn (down), or Admin Dn. Admin Dn indicates that the LSP is being taken down gracefully. |
| LSPname | Name of the LSP. |
| Setup priority | Value of the setup priority. |
| Hold priority | Determines the degree to which an LSP holds onto its session reservation after the LSP has been set up successfully. |
| CSPF usage | CSPF usage state in the rsvp session. |
| LSP Protection | Protects the traffic failures. |
| Label in | Incoming label for this LSP. |
| Label out | Outgoing label for this LSP. |
| Tspec rate | Sender's traffic specification, which describes the sender's traffic parameters. |
| Fspec rate | Fspec peak rate values. |
| Tunnel id | Tunnel address (destination port) for the session. |
| LSP id | Address of the LSP in the interface. |
| Ext-Tunnel Id | Session address for the ext-tunnel. |
| Down stream | Specify the dstream label for the bidirectional label-switched path (LSP). |
| Path refresh | Path messages are sent periodically to refresh path states. The refresh interval is controlled by a variable called the refresh time. |
| Resv lifetime | Number of seconds remaining in the lifetime of the reservation. |
| Retry count | Number of times sanity polling periodically checks for an error condition in the FPC. |
| | Interval sets the time for the messages in order to control the session. |
| RRO re-use as ERO | Enabling to re-use Record route as Explicit route for rsvp session. |
| Label Recording | Enabling to record the labels exchanged by all the peers. |
| Admin Groups | Resource affinities associated with the rsvp session. |
| Configured Path | Configured path name associated with the rsvp session. |
| Configured Explicit Route Detail | Configured explicit route with each hop information. |

Table 48. show rsvp session LSP-NAME output field (continued)

| Field | Description |
|-------------------------------|--|
| Session Explicit Route Detail | Established explicit route with each hop information. |
| Record route | Established rsvp path with each hop information. |
| Style | Reservation style associated with the rsvp session. |
| Traffic type | Traffic type associated with the rsvp session. |
| Minimum Path MTU | Path maximum transmission unit (MTU) discovery in the interface. |
| Last Recorded Error Code | Recorded error code for the last time service ran. |
| Last Recorded Error Value | No Recorded error value for the last time service ran. |

show rsvp session transit

Use this command to display session-related information for the transit or intermediate router.

Command Syntax

```
show rsvp session transit
show rsvp session transit detail
show rsvp session transit up
show rsvp session transit down
show rsvp session transit up detail
show rsvp session transit down detail
```

Parameters

up

Use this parameter to display sessions that are operational

down

Use this parameter to display sessions that are not operational

detail

Use this parameter to display detailed session-related information

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

Following are sample outputs from the command displaying detailed session information for the transit router.

```
#show rsvp session transit detail
Transit (Primary)
10.10.21.3
  From: 1.1.1.1, LSPstate: Up, LSPname: t1
  Setup priority: 5, Hold priority: 5
  LSP Protection: None
  Label in: 16, Label out: 3,
  Tspec rate: 10m, Fspec rate: 10m
  Tunnel Id: 1, LSP Id: 2, Ext-Tunnel Id: 1.1.1.1
  Downstream: 10.10.21.3, eth1 Upstream: 10.10.23.1, eth3
  Path refresh: 5 seconds (due in 6155 seconds)
  Path lifetime: 26 seconds (due in 25 seconds)
  Resv refresh: 5 seconds (due in 2533 seconds)
  Resv lifetime: 26 seconds (due in 25 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: Received Explicit Route Detail :
    10.10.23.2/32 strict
  Record route: 10.10.23.1 <self> 10.10.21.3
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type: ELSP_SIGNAL
```



```

CLASS      DSCP_value      EXP_value
af43       100110          7
DSTE Class Type Number: 0, Class Type name: default
#

```

[Table 49](#) explains the show command output fields.

Table 49. show rsvp session transit output field

| Field | Description |
|----------------------------------|---|
| Transit (Primary) | Transit RSVP sessions information in the interface. |
| From | Source (ingress switch) of the session. |
| LSP state | State of the LSP that is being handled by this RSVP session. It can be either Up, Dn (down), or Admin Dn. Admin Dn indicates that the LSP is being taken down gracefully. |
| LSP name | Name of the LSP. |
| Setup priority | Value of the setup priority. |
| Hold priority | Determines the degree to which an LSP holds onto its session reservation after the LSP has been set up successfully. |
| LSP Protection | Protects the traffic failures. |
| Label in | Incoming label for this LSP. |
| Label out | Outgoing label for this LSP. |
| Tspec rate | Sender's traffic specification, which describes the sender's traffic parameters. |
| Fspec rate | Fspec peak rate values. |
| Tunnel id | Tunnel address (destination port) for the session. |
| LSP id | Address of the LSP in the interface. |
| Ext-Tunnel Id | Session address for the ext-tunnel. |
| Down stream | Specify the dnstream label for the bidirectional label-switched path (LSP). |
| Path refresh | Path messages are sent periodically to refresh path states. The refresh interval is controlled by a variable called the refresh time. |
| Resv lifetime | Number of seconds remaining in the lifetime of the reservation. |
| RRO re-use as ERO | Enabling to re-use Record route as Explicit route for rsvp session. |
| Label Recording | Enabling to record the labels exchanged by all the peers. |
| Admin Groups | Resource affinities associated with the rsvp session. |
| Configured Explicit Route Detail | Configured path name associated with the rsvp session. |
| Record route | Recorded route for the session, taken from the record route object. Normally this value will be the same as that of explit route. Differences indicate that path rerouting has occurred, typically during fast reroute. |
| Style | Reservation style associated with the rsvp session. |
| Traffic type | Traffic type associated with the rsvp session. |

Table 49. show rsvp session transit output field (continued)

| Field | Description |
|------------------------|--|
| Minimum Path MTU | Path maximum transmission unit (MTU) discovery in the interface. |
| LSP Type | Type of LSP for Diffserv services(E-LSP or L-LSP). |
| CLASS | Name of the class which is associated with rsvp session. |
| DSCP_value | DSCP value of diff-serv class which is associated with rsvp session. |
| EXP_value | EXP value of diff-serv class which is associated with rsvp session. |
| DSTE Class Type Number | Diff-serv class type number associated with rsvp session. |
| Class Type name | Diff-serv class type name associated with rsvp session. |

show rsvp statistics

Use this command to display overall statistics of different type of RSVP control messages sent and received in a node.

Command Syntax

```
show rsvp statistics
```

Parameters

None

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp statistics
PacketType      Sent      Received
Path            627        501
PathErr         0          24
PathTear        1          27
Resv FF         30          9
Resv WF         0          0
Resv SE         646       583
Resv Err        0          0
ResvTear        0          0
ResvConf        0          0
Hello          330604    334461
Bundle         1006       866
Ack             50         14
SRefresh       34348    32424
Notify          0          0
```

show rsvp summary-refresh

Use this command to display RSVP summary refresh data.

Command Syntax

```
show rsvp summary-refresh
```

Parameters

None

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp summary-refresh:
Neighbor Addr      Tunnel ID  LSP ID      Ingress      Egress
```

[Table 50](#) explains the show command output fields.

Table 50. show rsvp trunk output field

| Field | Description |
|---------------|--|
| Neighbor Addr | Neighbor address on the primary address of the interface. |
| Tunnel ID | Tunnel identifier (destination port) for the RSVP session. |
| LSP ID | Address of the LSP in the interface. |
| Ingress | Information about ingress RSVP sessions. |
| Egress | Information about egress RSVP sessions. |

show rsvp trunk

Use this command to display information for a specific trunk or for all trunks.

Command Syntax

```
show rsvp trunk
show rsvp trunk NAME
show rsvp trunk detail
```

Parameters

NAME

Enter the name of a trunk

detail

Use this parameter to display detailed information for all trunks

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp trunk
Trunk Name      Trunk ID  Type      # Sess      Egress Address(es)
T1              101      P2P       1           4.4.4.4
T2              102      P2P       2           5.5.5.5
Total trunks configured: 3.
#
```

Following is a sample output from the command using the detail parameter.

```
#show rsvp trunk detail
Trunk name: T1, tunnel-id: 101
Type: P2P
Ext-tunnel-id: 1.1.1.1/32
Egress: 4.4.4.4/32
# of LSPs in trunk: 1
Mapped-routes: none

Trunk name: T2, tunnel-id: 102
Type: P2P
Ext-tunnel-id: 1.1.1.1/32
Egress: 5.5.5.5/32
# of LSPs in trunk: 2
Mapped-routes: none
```

[Table 51](#) explains the show command output fields.

Table 51. show rsvp trunk output field

| Field | Description |
|-------------------------|---|
| Trunk Name | Name of the trunk. |
| Trunk ID | Session address for the trunk. |
| Type | Trunk type in the rsvp session. |
| Sess | Number of sessions associated with rsvp trunk. |
| Egress | Information about egress RSVP sessions. |
| Total trunks configured | Number of configured trunk in the rsvp session. |
| Ext-tunnel-id | Extended Tunnel identifier (destination port) for the RSVP session. |
| Mapped-routes | Map the route of the interface. |

show rsvp trunk multi-sec-detail

Use this command to display secondary priority details specific to a trunk or for all trunks

Command Syntax

```
show rsvp trunk multi-sec-detail
show rsvp trunk NAME multi-sec-detail
```

Parameters

NAME

Enter the name of a trunk

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
# show rsvp trunk multi-sec-detail
Ingress (Secondary-Priority1)
5.5.5.5
  From: 6.6.6.6, LSPstate: Dn, LSPname: t2-Secondary-Priority-1
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: -, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2206, Ext-Tunnel Id: 6.6.6.6
  Last Recorded Error Code: Routing Problem (24)
  Last Recorded Error Value: No route available toward destination (5)
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Ingress (Secondary-Priority3)
5.5.5.5
  From: 6.6.6.6, LSPstate: Dn, LSPname: t2-Secondary-Priority-3
  Ingress FSM state: Idle
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: -, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2206, Ext-Tunnel Id: 6.6.6.6
  Last Recorded Error Code: Routing Problem (24)
  Last Recorded Error Value: No route available toward destination (5)
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
```

show rsvp trunk-id-table

Use this command to display summary information for the RSVP trunk-ID allocation table. The command provides block-level details including index ranges, usable and free indices, and the first available trunk-ID index.

It also displays the total number of used and free indices across all allocated blocks.

Command Syntax

```
show rsvp trunk-id-table
show rsvp trunk-id-table block-id <BLOCK-ID>
```

Parameters

BLOCK-ID

Specifies the block-id to display detailed bitmap information for a specific trunk-ID block, including per-bit allocation status.

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

Following is a sample output from the command using the detail parameter.

```
vpc2#show rsvp trunk-id-table
Bitmap Table: trunk-id, Range: 5001-65535, Current block: 3
+-----+-----+-----+-----+-----+
block_id  index_min  index_max  usable_indices  free_indices  first_free_index
+-----+-----+-----+-----+-----+
3         5001      5119      119             118           5002

Total - blocks: 1, used-indices: 1, free-indices: 118

vpc2#
```

Following is a sample output from the command using the BLOCK-ID parameter.

```
vpc2#show rsvp trunk-id-table block-id 3
Block ID           : 3
Bitmap Table       : trunk-id
Range              : 5001-5119
Usable indices     : 119
Free index count   : 118
First free index   : 5002
Flags              : NONE
Bitmap array (1 - used, 0 - free) :
[ 4992 ]          1 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 5023 ]
[ 5024 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 5055 ]
[ 5056 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 5087 ]
[ 5088 ] 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 [ 5119 ]
```



```
vpc2#
```

show rsvp version

Use this command to display the version of the RSVP daemon. Current RSVP version is 1.

Command Syntax

```
show rsvp version
```

Parameters

None

Command Mode

Exec and Privileged Exec modes

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show rsvp version
Resource ReSerVation Protocol, version 1. rfc2205
  RSVP protocol      = Enabled
  R(refresh timer)    = 30 seconds
  K(keep multiplier)  = 3
  Preemption          = Normal
#
```

[Table 52](#) explains the show command output fields.

Table 52. show rsvp version output field

| Field | Description |
|-------------------------------|--|
| Resource Reservation Protocol | RSVP software version. |
| RSVP protocol | Status of RSVP. |
| R (refresh timer) | Configured time interval used to generate periodic RSVP messages. |
| K (keep multiplier) | Number of RSVP messages that can be lost before an RSVP state is declared stale. |
| Preemption | Currently configured preemption capability. |

MULTI-PROTOCOL LABEL SWITCHING CONFIGURATION

| | |
|---|------|
| Understanding Label Space | 1106 |
| Overview | 1106 |
| Topology | 1106 |
| Configuration | 1106 |
| Validation | 1109 |
| Understanding MPLS TTL Processing | 1112 |
| Overview | 1112 |
| Topology | 1112 |
| Configuration | 1113 |
| Validation | 1115 |
| MPLS-TE Shared Risk Link Group | 1117 |
| Configure MPLS-TE Shared Risk Link Group | 1117 |
| RSVP SRLG with ISIS IGP | 1147 |
| Overview | 1147 |
| Characteristics | 1147 |
| Benefits | 1148 |
| SRLG with ISIS IGPConfiguration | 1148 |
| SRLG with ISIS IGP Commands | 1193 |
| Implementation Examples | 1194 |
| Troubleshooting | 1196 |
| Glossary | 1196 |
| MPLS DiffServ Configuration | 1198 |
| MPLS Diff-Serv Overview | 1198 |
| Terminology | 1199 |
| Remarking Configuration | 1201 |
| Configuration | 1201 |
| OSPF and LDP Configuration for PE1, P and PE2 | 1201 |
| Policing Configuration | 1219 |
| Configuration | 1219 |
| Validation | 1224 |
| MPLS Statistics Configuration | 1230 |
| Topology | 1230 |
| Configure LDP-LSP | 1230 |
| Configure RSVP-LSP | 1237 |
| Configure Static-LSP | 1238 |

| | |
|--|------|
| Validation | 1240 |
| MPLS Service Mapping Configuration | 1245 |
| Overview | 1245 |
| Topology | 1245 |
| Configuration | 1247 |
| Validation | 1262 |
| Mapping RSVP Tunnel Name to L2VPN Service | 1269 |
| Topology | 1269 |
| Configuring the Virtual Circuit (VC) | 1269 |
| Configure IP Address and OSPF on Routers | 1270 |
| Configure MPLS, RSVP, and LDP Targeted Peer on Routers | 1271 |
| Configure VC | 1273 |
| Bind Customer Interface to VC | 1274 |
| Configuring a MPLS Static Layer-2 VC | 1277 |
| Configure Dynamic VPLS | 1282 |
| LDP VPLS Service Mapping Configuration | 1283 |
| Configure Static VPLS | 1286 |
| MPLS BFD Configuration | 1289 |
| Topology | 1289 |
| Configuring MPLS BFD LSP over LDP | 1289 |
| Configuring MPLS BFD LSP over RSVP | 1293 |
| Configuring Static MPLS BFD | 1296 |
| Configuring PW VCCV BFD | 1297 |
| BGP Peer Groups for Address-Family L2VPN EVPN | 1301 |
| Topology | 1301 |
| Configuration | 1301 |
| Validation | 1304 |
| Point-to-Point Connections Over MPLS | 1308 |
| Overview | 1308 |
| Topology | 1308 |
| Configure IP Address and OSPF on Routers | 1308 |
| Configure MPLS, LDP, and LDP Targeted Peer on Routers | 1311 |
| Configure VC | 1313 |
| Bind Customer Interface to VC | 1314 |
| Validation | 1315 |
| L2VPN-L3VPN Stitching | 1318 |
| Overview | 1318 |
| Configuration | 1319 |
| New CLI | 1373 |
| Implementation Examples | 1374 |

| | |
|-----------------------|------|
| Troubleshooting | 1375 |
| Glossary | 1375 |

Understanding Label Space

This chapter contains configurations for Label Space. It also provides an overview of Label Space concepts.


Overview

The Label space refers to the scope of labels in a given LSR. It determines assignment and distribution of labels to a given peer. During data flow, it decides the key for looking up MPLS table and takes appropriate action based on the entry.

OcNOS supports **per-platform** label space, where a label must be unique for the entire platform. A label is interpreted the same way at all the interfaces. The FIB entry in the router does not contain incoming interface-related information. Thus the incoming traffic will be matched only with the label.

Topology

Figure 37. LDP Topology



Configuration

PE1

| | |
|--|---|
| PE1#configure terminal | Enter configure mode |
| PE1(config)#interface lo | Enter interface mode. |
| PE1(config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback address |
| PE1(config-if)#exit | Exit interface mode |
| PE1(config)#interface xe6 | Specify the interface (xe6)to be configured |
| PE1(config-if)#ip address 10.10.10.1/24 | Configure IP address for the interface |
| PE1(config-if)#no shutdown | Administratively bringing up the interface |
| PE1(config-if)#exit | Exit interface mode |
| PE1(config)#router ospf 100 | Configure the routing process and specify the |

| | |
|---|---|
| | Process ID (100) |
| PE1(config-router)#ospf router-id 1.1.1.1 | Configure OSPF router ID same as loopback IP address |
| PE1(config-router)#network 10.10.10.0/24 area 0 PE1(config-router)#network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID(0)with the interface |
| PE1(config-router)#exit | Exit configure mode |
| PE1(config)#router ldp | Enter router mode for LDP |
| PE1(config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1 |
| PE1(config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter "ipv6" if you are configuring an IPv6 interface. |
| PE1(config-router)# targeted-peer ipv4 3.3.3.3 | Configure targeted peer. |
| PE1(config-router-targeted-peer)#exit | Exit from router target peer and LDP mode |
| PE1(config-if)#exit | Exit interface mode |
| PE1(config)#interface xe6 | Specify the interface (xe6)to be configured |
| PE1(config-if)#label-switching | Enabling label switching capability on router |
| PE1(config-if)#enable-ldp ipv4 | Enabling ldp on interface |
| PE1(config-if)#commit | Commit the transaction |

P

| | |
|--|---|
| P#configure terminal | Enter configure mode. |
| P(config)#interface lo | Enter interface mode. |
| P(config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback address |
| P(config-if)#exit | Exit interface mode |
| P(config)#interface xe6 | Specify the interface (xe6)to be configured |
| P(config-if)#ip address 10.10.10.2/24 | Configure IP address for the interface |
| P(config-if)#no shutdown | Administratively bringing up the interface |
| P(config)#interface xe1 | Specify the interface (xe1)to be configured |
| P(config-if)#ip address 20.20.20.1/24 | Configure IP address for the interface |
| P(config)#router ospf 100 | Configure the routing process and specify the Process ID (100) |
| P(config-router)#ospf router-id 2.2.2.2 | Configure OSPF router ID same as loopback IP address |
| P(config-router)#network 10.10.10.0/24 area 0 P(config-router)#network 20.20.20.0/24 area 0 P(config-router)#network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0)with the interface |
| P(config-router)#exit | Exit router mode |

| | |
|---|---|
| P(config)#router ldp | Enter router mode for LDP |
| P(config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| P(config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter "ipv6" if you are configuring an IPv6 interface. |
| P(config-router)#exit | Exit router mode for LDP |
| P(config)#mpls min-label-value 1000 max-label-value 50000 label-space 0 | Configure the minimum label value and maximum label value to be used by Platform label space (Label space 0) |
| P(config)#interface xe6 | Specify the interface (xe6)to be configured |
| P(config-if)#label-switching | Enabling label switching capability on router |
| P(config-if)#enable-ldp ipv4 | Enabling ldp on interface |
| P(config-if)#exit | Exit interface mode |
| P(config)#interface xe1 | Specify the interface (xe1) to be configured |
| P(config-if)#label-switching | Enabling label switching capability on router |
| P(config-if)#enable-ldp ipv4 | Enabling ldp on interface |
| P(config-if)#commit | Commit the transaction |

PE2

| | |
|---|---|
| PE2#configure terminal | Enter configure mode. |
| PE2(config)#interface lo | Enter interface mode. |
| PE2(config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback address |
| PE2(config-if)#exit | Exit interface mode |
| PE2(config)#interface xe1 | Specify the interface (xe1)to be configured |
| PE2(config-if)#ip address 20.20.20.2/24 | Configure IP address for the interface |
| PE2(config-if)#no shutdown | Administratively bringing up the interface |
| PE2(config-if)#exit | Exit interface mode |
| PE2(config)#router ospf 100 | Configure the routing process and specify the Process ID (100) |
| PE2(config-router)#ospf router-id 3.3.3.3 | Configure OSPF router ID same as loopback IP address |
| PE2(config-router)#network 20.20.20.0/24 area 0 PE2(config-router)#network 3.3.3.3/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0)with the interface |
| PE2(config-router)#exit | Exit router mode |
| PE2(config)#router ldp | Enter router mode for LDP |
| PE2(config-router)#router-id 3.3.3.3 | Set the router ID to IP address 3.3.3.3 |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address to be used for a TCP |

| | |
|--|---|
| | session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter "ipv6" if you are configuring an IPv6 interface. |
| PE2(config-router)# targeted-peer ipv4 1.1.1.1 | Configure targeted peer |
| PE1(config-router-targeted-peer)#exit | Exit from router target peer and LDP mode |
| PE2(config-router)#exit | Exit router mode for LDP |
| PE2(config)#interface xe1 | Specify the interface (xe1) to be configured |
| PE2(config-if)#label-switching | Enabling label switching capability on router |
| PE2(config-if)#enable-ldp ipv4 | Enabling ldp on interface |
| PE2(config-if)#commit | Commit the transaction. |

Validation

```

PE1#show ldp
Router ID           : 1.1.1.1
LDP Version         : 1
Fast-reroute Per-prefix : Disabled
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode   : Liberal
Label Control Mode     : Independent
Instance Loop Detection : Off
Request Retry          : Off
Propagate Release      : Disabled
Graceful Restart       : Disabled
Hello Interval         : 5
Targeted Hello Interval : 15
Hold time              : 15
Targeted Hold time     : 45
Keepalive Interval     : 10
Keepalive Timeout      : 30
Request retry Timeout   : 5
Auto Targeted Hello Receipt : Disabled
Transport Address data :
  Label space 0        : 1.1.1.1 (in use)
Import BGP routes      : No

PE1#show ip ospf neighbor

```

```

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID   Pri   State           Dead Time   Address        Interface      Instance ID
2.2.2.2       1    Full/Backup     00:00:30    10.10.10.2     xe1            0

```

```

P#show ldp

Router ID           : 2.2.2.2
LDP Version         : 1
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode   : Liberal
Label Control Mode     : Independent
Instance Loop Detection : Off
Request Retry          : Off
Propagate Release      : Disabled

```

```

Graceful Restart      : Disabled
Hello Interval        : 5
Targeted Hello Interval : 15
Hold time             : 15
Targeted Hold time    : 45
Keepalive Interval    : 10
Keepalive Timeout     : 30
Request retry Timeout : 5
Auto Targeted Hello Receipt: Disabled
Transport Address data :
  Labelspace 0        : 2.2.2.2 (in use)
Import BGP routes     : No

```

P#show mpls label-space 0

```

Label range (min - max)      : 1000 - 50000

module-static
  Default range              : 1000 - 3839

module-srgb
  Default range (Usable)     : 3840 - 8739
  Default range (Allotted)   : 3840 - 8959

module-srlb
  Default range (Usable)     : 8960 - 9959
  Default range (Allotted)   : 8960 - 10239

module-rsvp
  Configured range           : N/A
  Current dynamic range      : N/A

module-ldp
  Configured range           : Not configured
  Current dynamic range      : 10880 - 11519

module-ldp-vc
  Configured range           : Not configured
  Current dynamic range      : 11520 - 12159

module-bgp
  Configured range           : N/A
  Current dynamic range      : N/A

```

P#show ip ospf neighbor

```

Total number of full neighbors: 2
OSPF process 100 VRF(default):

```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|--------------|-----|-------------|-----------|------------|-----------|-------------|
| 10.12.49.142 | 1 | Full/DR | 00:00:34 | 10.10.10.1 | xe1 | 0 |
| 10.12.49.158 | 1 | Full/Backup | 00:00:36 | 20.20.20.2 | xe6 | 0 |

P#show mpls ilm-table

```

Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
S - SNMP, L - LDP, R - RSVP, C - CRLDP
B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
O - OSPF/OSPF6 SR, i - ISIS_SR, k - SR CLI
P - SR Policy, U - unknown

```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|---------------|---------------|-------------|----------|-----------|---------|------|
| Intf/VRF | Nextthop | | pri | LSP-Type | | |
| L> 1.1.1.1/32 | | 2 | 10881 | 3 | N/A | eth3 |
| | Yes | LSP_DEFAULT | | | | |
| L> 3.3.3.3/32 | | 1 | 10880 | 3 | N/A | eth2 |
| | Yes | LSP_DEFAULT | | | | |

PE2#show ldp

```
Router ID           : 3.3.3.3
LDP Version         : 1
Fast-reroute Per-prefix : Disabled
Global Merge Capability : Merge Capable
Label Advertisement Mode : Downstream Unsolicited
Label Retention Mode   : Liberal
Label Control Mode     : Independent
Instance Loop Detection : Off
Request Retry         : Off
Propagate Release      : Disabled
Graceful Restart       : Disabled
Hello Interval        : 5
Targeted Hello Interval : 15
Hold time             : 15
Targeted Hold time     : 45
Keepalive Interval     : 10
Keepalive Timeout      : 30
Request retry Timeout   : 5
Auto Targeted Hello Receipt : Disabled
Transport Address data :
  Labelspace 0         : 3.3.3.3 (in use)
Import BGP routes      : No
```

```
PE2#show ip ospf neighbor
```

```
Total number of full neighbors: 1
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 2.2.2.2 | 1 | Full/Backup | 00:00:29 | 20.20.20.1 | xe1 | 0 |

Understanding MPLS TTL Processing

This chapter contains configurations for MPLS-TTL-Processing. It also provides an overview of MPLS-TTL-Processing concepts.

Overview

This feature performs 'Time To Live' (TTL) processing for Multi-Protocol Label Switching (MPLS) packets. The TTL processing is decided by the model chosen by you. This feature provides TTL processing of MPLS packets on ingress, egress, and intermediate routers. TTL processing is compliant with RFC 3443.


The details of TTL processing vary with the tunnel model that is configured for TTL processing. The incoming and outgoing TTL of the packet is determined by the configured tunnel model. Two Models are supported, pipe model and uniform model. Pipe model is default, where MPLS header TTL Value won't get propagated to IP header.

While pipe mode is the default, you can choose the uniform model with the [mpls lsp-model uniform \(page 1419\)](#) command.

For more about the uniform model and pipe models, see .

Topology

Figure 38. TTL Processing Topology



Configuration

PE1

| | |
|---|--|
| PE1#configure terminal | Enter configure mode |
| PE1(config)#interface lo | Enter interface mode. |
| PE1(config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback address |
| PE1(config-if)#exit | Exit interface mode |
| PE1(config)#interface xe6 | Specify the interface (xe6)to be configured |
| PE1(config-if)#ip address 10.10.10.1/24 | Configure IP address for the interface |
| PE1(config-if)#no shutdown | Administratively bringing up the interface |
| PE1(config-if)#exit | Exit interface mode |
| PE1(config)#router ospf 100 | Configure the routing process and specify the Process ID (100) |
| PE1(config-router)#ospf router-id 1.1.1.1 | Configure ospf router id same as loopback ip address. |
| PE1(config-router)#network 10.10.10.0/24 area 0 PE1(config-router)#network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID(0)with the interface |
| PE1(config-router)#exit | Exit configure mode |
| PE1(config)#router ldp | Enter router mode for LDP |
| PE1(config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1 |
| PE1(config-router)#transport-address ipv4 1.1.1.1 0 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. Note: It is preferable to use the loopback address as the transport address. |
| PE1(config-router)#targeted-peer ipv4 3.3.3.3 | Configure targeted peer. |
| PE1(config-router)#exit | Exit router mode for LDP |
| PE1(config)#interface xe6 | Specify the interface (xe6)to be configured |
| PE1(config-if)#label-switching | Enabling label switching capability on router |
| PE1(config-if)#enable-ldp ipv4 | Enabling ldp on interface |
| PE1(config-if)#commit | Commit the transaction |

P

| | |
|--|---|
| P#configure terminal | Enter configure mode. |
| P(config)#interface lo | Enter interface mode. |
| P(config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback address |

| | |
|--|--|
| P(config-if)#exit | Exit interface mode |
| P(config)#interface xe6 | Specify the interface (xe6)to be configured |
| P(config-if)#ip address 10.10.10.2/24 | Configure IP address for the interface |
| P(config-if)#no shutdown | Administratively bringing up the interface |
| P(config)#interface xe1 | Specify the interface (xe1)to be configured |
| P(config-if)#ip address 20.20.20.1/24 | Configure IP address for the interface |
| P(config)#router ospf 100 | Configure the routing process and specify the Process ID (100) |
| P(config-router)#ospf router-id 2.2.2.2 | Configure OSPF router id same as loopback IP address. |
| P(config-router)#network 10.10.10.0/24 area 0 P(config-router)#network 20.20.20.0/24 area 0 P(config-router)#network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0)with the interface |
| P(config-router)#exit | Exit router mode |
| P(config)#router ldp | Enter router mode for LDP |
| P(config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| P(config-router)#transport-address ipv4 2.2.2.2 0 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. Note: It is preferable to use the loopback address as the transport address. |
| P(config-router)#exit | Exit router mode for LDP |
| P(config)#interface xe6 | Specify the interface (xe6)to be configured |
| P(config-if)#label-switching | Enabling label switching capability on router |
| P(config-if)#enable-ldp ipv4 | Enabling ldp on interface |
| P(config)#interface xe1 | Specify the interface (xe1)to be configured |
| P(config-if)#label-switching | Enabling label switching capability on router |
| P(config-if)#enable-ldp ipv4 | Enabling ldp on interface |
| P(config-if)#commit | Commit the transaction |

PE2

| | |
|--|---|
| PE2#configure terminal | Enter configure mode. |
| PE2(config)#interface lo | Enter interface mode. |
| PE2(config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback address |
| PE2(config-if)#exit | Exit interface mode |
| PE2(config)#interface xe1 | Specify the interface (xe1)to be configured |
| PE2(config-if)#ip address 20.20.20.2/24 | Configure IP address for the interface |
| PE2(config-if)#no shutdown | Administratively bringing up the interface |

| | |
|---|--|
| PE2(config-if)#exit | Exit interface mode |
| PE2(config)#router ospf 100 | Configure the routing process and specify the Process ID (100) |
| PE2(config-router)#ospf router-id 3.3.3.3 | Configure OSPF router ID same as loopback IP address. |
| PE2(config-router)#network 20.20.20.0/24 area 0 PE2(config-router)#network 3.3.3.3/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0)with the interface |
| PE2(config)#router ldp | Enter router mode for LDP |
| PE2(config-router)#router-id 3.3.3.3 | Set the router ID to IP address 3.3.3.3 |
| PE2(config-router)#transport-address ipv4 3.3.3.3 0 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. Note: It is preferable to use the loopback address as the transport address. |
| PE2(config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| PE2 (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| PE2(config-router)#exit | Exit router mode for LDP |
| PE2(config)#interface xe1 | Specify the interface (xe1)to be configured |
| PE2(config-if)#label-switching | Enabling label switching capability on router |
| PE2(config-if)#enable-ldp ipv4 | Enabling ldp on interface |
| PE2(config-if)#commit | Commit the transaction |

Validation

```

PE1#show ip ospf neighbor
Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface      Instance ID
2.2.2.2        1     Full/DR          00:00:31    10.10.10.2     xe6            0

PE1#show ldp session
Peer IP Address      IF Name    My Role    State        KeepAlive  UpTime
3.3.3.3              xe6        Passive    OPERATIONAL  30         00:03:03
2.2.2.2              xe6        Passive    OPERATIONAL  30         00:03:46

P#show ip ospf neighbor
Total number of full neighbors: 2
OSPF process 100 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface      Instance ID
1.1.1.1        1     Full/Backup      00:00:37    10.10.10.1     xe6            0
3.3.3.3        1     Full/DR          00:00:34    20.20.20.2     xe1            0

P#show ldp session
Peer IP Address      IF Name    My Role    State        KeepAlive  UpTime
3.3.3.3              xe1        Passive    OPERATIONAL  30         00:06:21
1.1.1.1              xe6        Active     OPERATIONAL  30         00:06:39

```

```
PE2#show ip ospf neighbor
```

```
Total number of full neighbors: 1
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 2.2.2.2 | 1 | Full/Backup | 00:00:37 | 20.20.20.1 | xe1 | 0 |

```
PE2#show ldp session
```

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 1.1.1.1 | xe1 | Active | OPERATIONAL | 30 | 00:06:07 |
| 2.2.2.2 | xe1 | Active | OPERATIONAL | 30 | 00:06:33 |

```
PE1#ping 3.3.3.3
```

```
Press CTRL+C to exit
```

```
PING 3.3.3.3 (3.3.3.3) 56(84) bytes of data.
```

```
64 bytes from 3.3.3.3: icmp_seq=1 ttl=63 time=2.17 ms
```

```
64 bytes from 3.3.3.3: icmp_seq=2 ttl=63 time=2.26 ms
```

```
64 bytes from 3.3.3.3: icmp_seq=3 ttl=63 time=2.11 ms
```

```
64 bytes from 3.3.3.3: icmp_seq=4 ttl=63 time=1.91 ms
```

```
--- 3.3.3.3 ping statistics ---
```

```
4 packets transmitted, 4 received, 0% packet loss, time 8ms
```

```
rtt min/avg/max/mdev = 1.912/2.112/2.261/0.127 ms
```

```
PE1#traceroute 2.2.2.2
```

```
traceroute to 2.2.2.2 (2.2.2.2), 30 hops max, 60 byte packets
```

```
1 2.2.2.2 (2.2.2.2) 1.172 ms 0.918 ms 0.983 ms
```

```
PE1#traceroute 3.3.3.3
```

```
traceroute to 3.3.3.3 (3.3.3.3), 30 hops max, 60 byte packets
```

```
1 * * *
```

```
2 3.3.3.3 (3.3.3.3) 5.440 ms 5.215 ms 5.305 ms
```


MPLS-TE Shared Risk Link Group

Shared Risk Link Groups (SRLG) is a feature which allows the user to establish a backup secondary LSP (label switched path) path or a FRR (fast-reroute) LSP path which is disjoint from the path of the primary LSP. Links which are members of the same SRLG represent resources which share the same risk. For example, fiber links sharing the same conduit or multiple wavelengths sharing the same fiber.

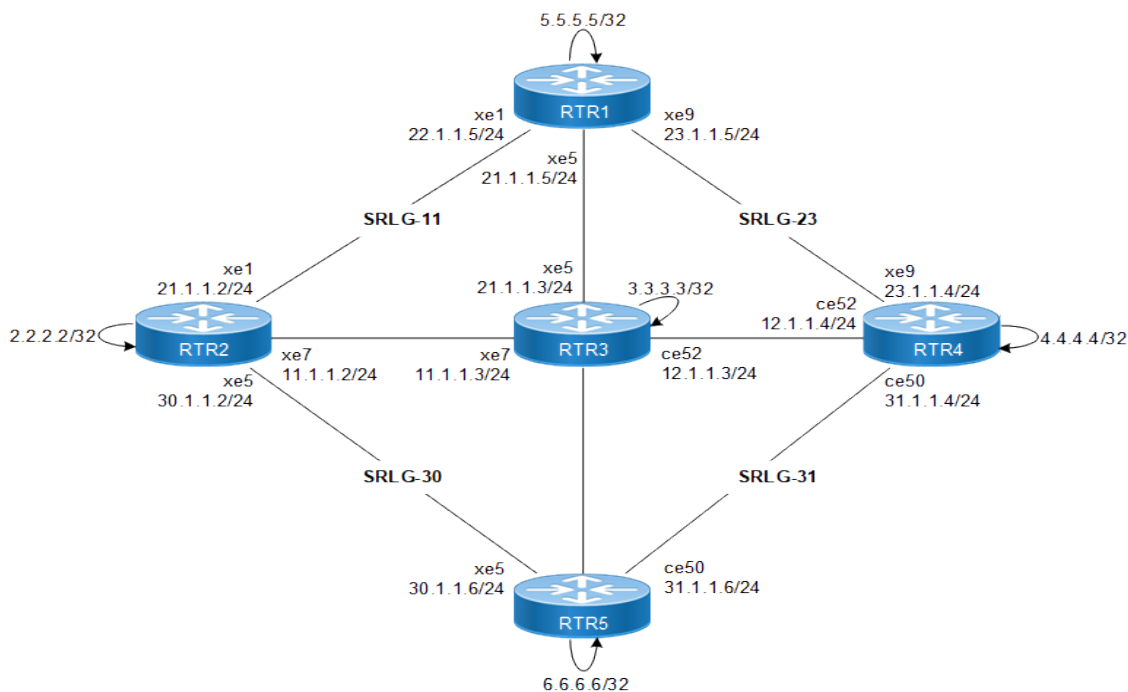
SRLG groups are used to determine which links belong to the same SRLG. To advertise SRLG, the information is part of the IGP TE parameters in an opaque LSA (link state advertisement). It is advertised in a SRLG sub-TLV (type 16) of the existing Link TLV in OSPF.

The SRLG of a path in a label-switched path (LSP) is the set of SRLGs for all the links in the path. When computing the secondary path for an LSP, it is preferable to find a path such that the secondary and primary paths do not have any links in common in case the SRLGs for the primary and secondary paths are disjoint. This ensures that a single point of failure on a particular link does not bring down both the primary and secondary paths in the LSP.

Backup tunnel path selection with SRLG is done so that a backup tunnel can avoid using links that are in the same SRLG as the interfaces it is protecting. Otherwise, when the protected link fails the backup tunnel fails too.

Configure MPLS-TE Shared Risk Link Group

Figure 39. Topology for SRLG



RTR2

Loopback Interface

```
#configure terminal
```

```
Enter configure mode
```

| | |
|--|---|
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address on loopback interface |
| (config-if)#commit | Commit the transaction |

MPLS Interface

| | |
|------------------------------------|---|
| (config)#interface xe1 | Enter the Interface mode |
| (config-if)#ip address 22.1.1.2/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface xe5 | Enter the Interface mode |
| (config-if)#ip address 30.1.1.2/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface xe7 | Enter the Interface mode |
| (config-if)#ip address 11.1.1.2/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#commit | Commit the transaction. |

OSPF

| | |
|--|------------------------------------|
| (config)#router ospf 1 | Enter the Router OSPF mode |
| (config-router)#router-id 2.2.2.2 | Configure OSPF router ID |
| (config-router)#network 2.2.2.2/32 area 0 | Advertise loopback address in OSPF |
| (config-router)#network 11.1.1.0/24 area 0 | Advertise network address in OSPF |
| (config-router)#network 22.1.1.0/24 area 0 | Advertise network address in OSPF |
| (config-router)#network 30.1.1.0/24 area 0 | Advertise network address in OSPF |
| (config-if)#commit | Commit the transaction |

Global RSVP

| | |
|----------------------|----------------------------|
| (config)#router rsvp | Enter the Router RSVP mode |
| (config-if)#commit | Commit the transaction |

RSVP Interface

| | |
|-------------------------|---|
| (config)#interface xe1 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable RSVP on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface xe5 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable RSVP on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface xe7 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable rsvp on the interface |
| (config-if)#commit | Commit the transaction |

Global SRLG (Path Type: Forced)

| | |
|-------------------------------------|---------------------------------|
| (config)#router rsvp | Enter the Router RSVP mode. |
| config-router)#srlg-disjoint forced | Configure SRLG path type forced |
| (config-router)#commit | Commit the transaction. |



Note: Disjoint type “preferred” can be configured globally using the command `srlg-disjoint preferred`.

SRLG Interface

| | |
|--------------------------------------|--|
| (config)#interface xe1 | Enter the interface mode |
| (config-if)#mpls traffic-eng srlg 11 | Configure SRLG value in the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface xe7 | Enter the Interface mode |
| (config-if)#mpls traffic-eng srlg 11 | Configure SRLG value in the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface xe5 | Enter the Interface mode |
| (config-if)#mpls traffic-eng srlg 30 | Configure SRLG value in the interface |
| (config-if)#commit | Commit the transaction |

MPLS RSVP Path

| | |
|-------------------------------|-----------------------------------|
| (config)#rsvp-path path1 | Configure MPLS-path |
| (config-path)#11.1.1.3 strict | Configure first hop for the path |
| (config-path)#12.1.1.4 strict | Configure second hop for the path |

| | |
|-----------------------------------|--|
| (config-path)#exit | Exit mpls-path mode and return to Configure mode |
| (config)#rsvp-path sec-path1 | Configure MPLS-path |
| (config-path)#4.4.4.4 loose | Configure loose path |
| (config-path)#exit | Exit mpls-path mode and return to Configure mode |
| (config-path)#rsvp-path sec-path2 | Configure MPLS-path |
| (config-path)#4.4.4.4 loose | Configure loose path |
| (config-if)#commit | Commit the transaction |

RSVP-TE Tunnel with Secondary Protection and SRLG Disjoint Type Forced:

| | |
|--|--|
| #configure terminal | Enter configure mode |
| (config)#rsvp-trunk lsp1 | Configure RSVP trunk |
| (config-trunk)#primary path path1 | Configure primary path |
| (config-trunk)#secondary-priority 1 path sec-path1 | Configure secondary path |
| (config-trunk)#secondary-priority 1 srlg-disjoint forced | Configure SRLG for secondary path |
| (config-trunk)#secondary-priority 2 path sec-path2 | Configure secondary path |
| (config-trunk)#secondary-priority 2 srlg-disjoint forced | Configure SRLG for secondary path |
| (config-trunk)#to 4.4.4.4 | Configure RSVP trunk destination address |
| (config-if)#commit | Commit the transaction |



Note: Disjoint type preferred can be configured in rsvp-trunk secondary using the command "secondary-priority 2 srlg-disjoint forced.

RSVP-TE Tunnel with Fast Reroute Protection One-to-One

| | |
|---|--|
| #configure terminal | Enter configure mode |
| (config)#rsvp-trunk lsp2 | Configure RSVP trunk |
| (config-trunk)#primary fast-reroute protection one-to-one | Configure fast-reroute protection |
| (config-trunk)#primary path path1 | Configure primary path |
| (config-trunk)#to 4.4.4.4 | Configure RSVP trunk destination address |
| (config-if)#commit | Commit the transaction |

RSVP-TE Tunnel with Fast Reroute Protection Facility Backup

| | |
|---|--|
| #configure terminal | Enter configure mode |
| (config)#rsvp-bypass bypass1 | Configure rsvp bypass tunnel |
| (config-bypass)#path path1 | Configure path for bypass tunnel |
| (config-bypass)#to 4.4.4.4 | Configure destination address for tunnel |
| (config-bypass)#exit | Exit rsvp-bypass tunnel mode |
| (config)#rsvp-trunk lsp3 | Configure rsvp trunk |
| (config-trunk)#primary fast-reroute protection facility | Configure fast-reroute protection facility |
| (config-trunk)#primary path sec-path1 | Configure path |
| (config-trunk)#to 4.4.4.4 | Configure the destination of rsvp trunk |
| (config-if)#commit | Commit the transaction |

RTR3**Loopback Interface**

| | |
|---|---|
| #configure terminal | Enter configure mode |
| config#interface lo | Enter the Interface mode for loopback interface |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure IP address on loopback interface |
| (config-if)#commit | Commit the transaction |

MPLS Interfaces

| | |
|------------------------------------|---|
| (config)#interface xe5 | Enter the Interface mode |
| (config-if)#ip address 21.1.1.3/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface xe7 | Enter the Interface mode |
| (config-if)#ip address 11.1.1.3/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface ce52 | Enter the Interface mode |
| (config-if)#ip address 12.1.1.3/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#commit | Commit the transaction. |

OSPF

| | |
|--|------------------------------------|
| (config)#router ospf 1 | Enter the Router OSPF mode |
| (config-router)#router-id 3.3.3.3 | Configure OSPF router id |
| (config-router)#network 3.3.3.3/32 area 0 | Advertise loopback address in OSPF |
| (config-router)#network 11.1.1.0/24 area 0 | Advertise network address in OSPF |
| (config-router)#network 12.1.1.0/24 area 0 | Advertise network address in OSPF |
| (config-router)#network 21.1.1.0/24 area 0 | Advertise network address in OSPF |
| (config-if)#commit | Commit the transaction |

Global RSVP

| | |
|------------------------|----------------------------|
| (config)#router rsvp | Enter the router RSVP mode |
| (config-router)#commit | Commit the transaction |

RSVP Interface

| | |
|-------------------------|--|
| (config)#interface xe5 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable RSVP on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface xe7 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable rsvp on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface ce52 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable rsvp on the interface |
| (config-if)#commit | Commit the transaction |

Global SRLG (Path Type: Forced)

| | |
|--------------------------------------|---------------------------------|
| (config)#router rsvp | Enter the Router RSVP mode |
| (config-router)#srlg-disjoint forced | Configure SRLG path type forced |
| (config-router)#commit | Commit the transaction |

When you use the one-to-one or facility backup protection, you must configure the `srlg-disjoint` type on all routers. When you use the secondary protection type, you only need to configure the `srlg-disjoint` type at the ingress node.de.

SRLG Interface

| | |
|--------------------------------------|---|
| (config)#interface xe5 | Enter the Interface mode |
| (config-if)#mpls traffic-eng srlg 21 | Configure SRLG value in the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface xe7 | Enter the Interface mode |
| (config-if)#mpls traffic-eng srlg 11 | Configure SRLG value in the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface ce52 | Enter the Interface mode |
| (config-if)#mpls traffic-eng srlg 12 | Configure SRLG value in the interface |
| (config-if)#commit | Commit the transaction |

RTR4

Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter the Interface mode for loopback interface. |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure IP address on the loopback address. |
| (config-if)#commit | Commit the transaction. |

MPLS Interface

| | |
|------------------------------------|---|
| (config)#interface xe9 | Enter the Interface mode |
| (config-if)#ip address 23.1.1.4/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface ce52 | Enter the Interface mode |
| (config-if)#ip address 12.1.1.4/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface ce50 | Enter the Interface mode |
| (config-if)#ip address 31.1.1.4/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#commit | Commit the transaction. |

OSPF

| | |
|--|-------------------------------------|
| (config)#router ospf 1 | Enter Router OSPF mode. |
| (config-router)#router-id 4.4.4.4 | Configure OSPF router id |
| (config-router)#network 4.4.4.4/32 area 0 | Advertise loopback address in OSPF. |
| (config-router)#network 12.1.1.0/24 area 0 | Advertise network address in OSPF |
| (config-router)#network 23.1.1.0/24 area 0 | Advertise network address in OSPF |
| (config-router)#network 31.1.1.0/24 area 0 | Advertise network address in OSPF |
| (config-router)#commit | Commit the transaction |

Global RSVP

| | |
|------------------------|----------------------------|
| (config)#router rsvp | Enter the Router RSVP mode |
| (config-router)#commit | Commit the transaction |

RSVP Interface

| | |
|-------------------------|--|
| (config)#interface xe9 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable RSVP on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface ce50 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable RSVP on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface ce52 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable RSVP on the interface |
| (config-if)#commit | Commit the transaction |

Global SRLG (Path Type: Forced)

| | |
|--------------------------------------|---------------------------------|
| (config)#router rsvp | Enter the Router RSVP mode |
| (config-router)#srlg-disjoint forced | Configure SRLG path type forced |
| (config-router)#commit | Commit the transaction |

SRLG Interface

| | |
|-------------------------|--------------------------|
| (config)#interface ce50 | Enter the Interface mode |
|-------------------------|--------------------------|

| | |
|--------------------------------------|---|
| (config-if)#mpls traffic-eng srlg 31 | Configure SRLG value in the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface ce52 | Enter the Interface mode |
| (config-if)#mpls traffic-eng srlg 12 | Configure SRLG value in the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface xe9 | Enter the Interface mode |
| (config-if)#mpls traffic-eng srlg 23 | Configure SRLG value in the interface |
| (config-if)#commit | Commit the transaction |

RTR5

Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter the Interface mode for loopback interface. |
| (config-if)#ip address 5.5.5.5/32 secondary | Configure IP address on the loopback address. |
| (config-if)#commit | Commit the transaction. |

MPLS Interface

| | |
|------------------------------------|---|
| (config)#interface xe1 | Enter the Interface mode |
| (config-if)#ip address 22.1.1.5/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface xe9 | Enter the Interface mode |
| (config-if)#ip address 23.1.1.5/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |
| (config)#interface xe5 | Enter the Interface mode |
| (config-if)#ip address 21.1.1.5/24 | Configure ip address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#commit | Commit the transaction. |

OSPF

| | |
|------------------------|-------------------------|
| (config)#router ospf 1 | Enter Router OSPF mode. |
|------------------------|-------------------------|

| | |
|---|-------------------------------------|
| <code>(config-router)#router-id 5.5.5.5</code> | Configure OSPF router id |
| <code>(config-router)#network 5.5.5.5/32 area 0</code> | Advertise loopback address in OSPF. |
| <code>(config-router)#network 21.1.1.0/24 area 0</code> | Advertise network address in OSPF |
| <code>(config-router)#network 22.1.1.0/24 area 0</code> | Advertise network address in OSPF |
| <code>(config-router)#network 23.1.1.0/24 area 0</code> | Advertise network address in OSPF |
| <code>(config-router)#commit</code> | Commit the transaction |

Global RSVP

| | |
|-----------------------------------|----------------------------|
| <code>(config)#router rsvp</code> | Enter the Router RSVP mode |
| <code>(config)#commit</code> | Commit the transaction |

RSVP Interface

| | |
|--------------------------------------|--|
| <code>(config)#interface xe1</code> | Enter the Interface mode |
| <code>(config-if)#enable-rsvp</code> | Enable RSVP on the interface |
| <code>(config-if)#exit</code> | Exit Interface mode and return to Configure mode |
| <code>(config)#interface xe5</code> | Enter the Interface mode |
| <code>(config-if)#enable-rsvp</code> | Enable RSVP on the interface |
| <code>(config-if)#exit</code> | Exit Interface mode and return to Configure mode |
| <code>(config)#interface xe9</code> | Enter the Interface mode |
| <code>(config-if)#enable-rsvp</code> | Enable RSVP on the interface |
| <code>(config-if)#commit</code> | Commit the transaction |

Global SRLG (Path Type: Forced)

| | |
|---|---------------------------------|
| <code>(config)#router rsvp</code> | Enter the Router RSVP mode. |
| <code>(config-router)#srlg-disjoint forced</code> | Configure SRLG path type forced |
| <code>(config-router)#commit</code> | Commit the transaction. |

SRLG Interface

| | |
|---|---------------------------------------|
| <code>(config)#interface xe1</code> | Enter the Interface mode |
| <code>(config-if)#mpls traffic-eng srlg 11</code> | Configure SRLG value in the interface |

| | |
|---|--|
| <code>(config-if)#exit</code> | Exit Interface mode and return to Configure mode |
| <code>(config)#interface xe5</code> | Enter the Interface mode |
| <code>(config-if)#mpls traffic-eng srlg 21</code> | Configure SRLG value in the interface |
| <code>(config-if)#exit</code> | Exit Interface mode and return to Configure mode |
| <code>(config)#interface xe9</code> | Enter the Interface mode |
| <code>(config-if)#mpls traffic-eng srlg 23</code> | Configure SRLG value in the interface |
| <code>(config-if)#commit</code> | Commit the transaction |

RTR6

Loopback Interface

| | |
|--|--|
| <code>#configure terminal</code> | Enter configure mode |
| <code>(config)#interface lo</code> | Enter the Interface mode for loopback interface. |
| <code>(config-if)#ip address 6.6.6.6/32 secondary</code> | Configure IP address on the loopback address. |
| <code>(config-if)#commit</code> | Commit the transaction. |

MPLS Interface

| | |
|---|---|
| <code>(config)#interface xe5</code> | Enter the Interface mode |
| <code>(config-if)#ip address 30.1.1.6/24</code> | Configure ip address on the interface |
| <code>(config-if)#label-switching</code> | Enable label switching on the interface |
| <code>(config-if)#exit</code> | Exit Interface mode and return to Configure mode. |
| <code>(config)#interface ce50</code> | Enter the Interface mode |
| <code>(config-if)#ip address 31.1.1.6/24</code> | Configure ip address on the interface |
| <code>(config-if)#label-switching</code> | Enable label switching on the interface |
| <code>(config-if)#commit</code> | Commit the transaction |

OSPF

| | |
|--|-------------------------------------|
| <code>(config)#router ospf 1</code> | Enter Router OSPF mode |
| <code>(config-router)#router-id 6.6.6.6</code> | Configure OSPF router ID |
| <code>(config-router)#network 6.6.6.6/32 area 0</code> | Advertise loopback address in OSPF. |
| <code>(config-router)#network 30.1.1.0/24 area</code> | Advertise network address in OSPF. |

| | |
|--|------------------------------------|
| 0 | |
| (config-router)#network 31.1.1.0/24 area 0 | Advertise network address in OSPF. |
| (config-router)#commit | Commit the transaction. |

Global RSVP

| | |
|------------------------|-----------------------------|
| (config)#router rsvp | Enter the Router RSVP mode. |
| (config-router)#commit | Commit the transaction. |

RSVP Interface

| | |
|-------------------------|--|
| (config)#interface xe5 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable RSVP on the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface ce50 | Enter the Interface mode |
| (config-if)#enable-rsvp | Enable RSVP on the interface |
| (config-if)#commit | Commit the transaction |

Global SRLG (Path Type: Forced)

| | |
|--------------------------------------|---------------------------------|
| (config)#router rsvp | Enter the Router RSVP mode |
| (config-router)#srlg-disjoint forced | Configure SRLG path type forced |
| (config-router)#commit | Commit the transaction. |

SRLG Interface

| | |
|--------------------------------------|--|
| (config)#interface xe5 | Enter the Interface mode |
| (config-if)#mpls traffic-eng srlg 30 | Configure SRLG value in the interface |
| (config-if)#exit | Exit Interface mode and return to Configure mode |
| (config)#interface ce50 | Enter the Interface mode |
| (config-if)#mpls traffic-eng srlg 31 | Configure SRLG value in the interface |
| (config-if)#commit | Commit the transaction |

Validation

R1 Configuration (Ingress)

OSPF:

```
RTR2#show ip ospf neighbor
Total number of full neighbors: 3
OSPF process 1 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|----------|-----------|-------------|
| 3.3.3.3 | 1 | Full/DR | 00:00:29 | 11.1.1.3 | xe7 | 0 |
| 5.5.5.5 | 1 | Full/DR | 00:00:37 | 22.1.1.5 | xe1 | 0 |
| 6.6.6.6 | 1 | Full/DR | 00:00:35 | 30.1.1.6 | xe5 | 0 |

RSVP:

```
RTR2#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
```

Ingress RSVP:

| To | From | Type | LSPName | State | | | | |
|---------|-------|----------|------------------|---------------------------|----|----------|--------|---|
| Uptime | Rt | Style | Labelin Labelout | DSType | | | | |
| 4.4.4.4 | 24960 | 2.2.2.2 | PRI | lsp1-Primary | UP | 01:14:21 | 1 1 SE | - |
| 4.4.4.4 | 24320 | 2.2.2.2 | SEC | lsp1-Secondary-Priority-1 | UP | 01:14:21 | 1 1 SE | - |
| 4.4.4.4 | 24961 | 2.2.2.2 | PRI | lsp2-Primary | UP | 01:03:18 | 1 1 SE | - |
| 4.4.4.4 | 24321 | 30.1.1.2 | DTR | lsp2-Detour | UP | 01:03:18 | 1 1 SE | - |
| 4.4.4.4 | 24962 | 2.2.2.2 | BPS | bypass1-Bypass | UP | 00:58:02 | 1 1 SE | - |
| 4.4.4.4 | 24321 | 2.2.2.2 | PRI | lsp3-Primary | UP | 00:57:26 | 1 1 SE | - |

Total 6 displayed, Up 6, Down 0.

RTR2#show rsvp session detail

Ingress (Primary)

```
4.4.4.4
  From: 2.2.2.2, LSPstate: Up, LSPname: lsp1-Primary
  Ingress FSM state: Operational
  Establishment Time: 0s 16ms
  SRLG configured in primary path are: 11 12
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
  IGP-Shortcut: Disabled, LSP metric: 201
  LSP Protection: None
  Label in: -, Label out: 24960,
  Tspec rate: 0, Espec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 2.2.2.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 11.1.1.3, xe7
  Path refresh: 30 seconds (RR enabled) (due in 25417 seconds)
  Resv lifetime: 157 seconds (due in 133 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: path1 (in use)
  Configured Explicit Route Detail :
    11.1.1.3/32 strict
```

```

12.1.1.4/32 strict
Session Explicit Route Detail :
11.1.1.3/32 strict
12.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
<self>
11.1.1.3
12.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Ingress (Secondary-Priority1)
4.4.4.4
  From: 2.2.2.2, LSPstate: Up, LSPname: lsp1-Secondary-Priority-1
  Ingress FSM state: Operational
  Establishment Time: 0s 38ms
  SRLG configured in backup path are: 30 31
  SRLG-disjoint Configured: Forced
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
  IGP-Shortcut: Disabled, LSP metric: 101
  LSP Protection: None
  Label in: -, Label out: 24320,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2204, Ext-Tunnel Id: 2.2.2.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 30.1.1.6, xe5
  Path refresh: 30 seconds (RR enabled) (due in 25416 seconds)
  Resv lifetime: 157 seconds (due in 127 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: sec-path1 (in use)
  Configured Explicit Route Detail :
    4.4.4.4/32 loose
  Session Explicit Route Detail :
    30.1.1.6/32 strict
    31.1.1.4/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  30.1.1.6
  31.1.1.4
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Ingress (Primary)
4.4.4.4
  From: 2.2.2.2, LSPstate: Up, LSPname: lsp2-Primary
  Ingress FSM state: Operational
  Establishment Time: 0s 10ms
  SRLG configured in primary path are: 11 12

```

```

Setup priority: 7, Hold priority: 0
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
IGP-Shortcut: Disabled, LSP metric: 201
LSP Protection: one-to-one
Fast-Reroute bandwidth : 0
Protection type desired: Link
Fast-Reroute Hop limit: 255
Fast-Reroute Setup priority: 7, Hold priority: 0
Label in: -, Label out: 24961,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5002, LSP Id: 2206, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 11.1.1.3, xe7
Path refresh: 30 seconds (RR enabled) (due in 26088 seconds)
Resv lifetime: 157 seconds (due in 133 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: path1 (in use)
Configured Explicit Route Detail :
  11.1.1.3/32 strict
  12.1.1.4/32 strict
Session Explicit Route Detail :
  11.1.1.3/32 strict
  12.1.1.4/32 strict
Record route:
LP = 1 -> PLR's Downstream link is protected      PU = 1 -> Protection is in use on PLR
NP = 1 -> PLR's Downstream neighbor is protected  BP = 1 -> BW protection available at PLR
-----
IP Address      Label      (LP, PU, NP, BP)
-----
<self>
11.1.1.3        ( 1,  0,  0,  0)
12.1.1.4        ( 0,  0,  0,  0)
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Ingress (Detour)
4.4.4.4
From: 30.1.1.2, LSPstate: Up, LSPname: lsp2-Detour
Ingress FSM state: Operational
Establishment Time: 0s 9ms
SRLG configured in backup path are: 30 31
SRLG-disjoint Configured: Forced
Setup priority: 7, Hold priority: 0
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
IGP-Shortcut: Disabled, LSP metric: 101
LSP Protection: None
Label in: -, Label out: 24321,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5002, LSP Id: 2206, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 30.1.1.6, xe5
Path refresh: 30 seconds (RR enabled) (due in 26073 seconds)
Resv lifetime: 157 seconds (due in 127 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none

```

```

Exclude path detail:
  Exclude Link: 11.1.1.3
Configured Path: none
Session Explicit Route Detail :
  30.1.1.6/32 strict
  31.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
<self>
30.1.1.6
31.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Ingress (Bypass)
4.4.4.4
  From: 2.2.2.2, LSPstate: Up, LSPname: bypass1-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 10ms
  SRLG configured in primary path are: 11 12
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
  IGP-Shortcut: Disabled, LSP metric: 201
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
    Label in: -, Label out: 24962,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5003, LSP Id: 2207, Ext-Tunnel Id: 2.2.2.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 11.1.1.3, xe7
  Path refresh: 30 seconds (RR enabled) (due in 26388 seconds)
  Resv lifetime: 157 seconds (due in 133 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: path1 (in use)
  Configured Explicit Route Detail :
    11.1.1.3/32 strict
    12.1.1.4/32 strict
  Session Explicit Route Detail :
    11.1.1.3/32 strict
    12.1.1.4/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  11.1.1.3
  12.1.1.4
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
  Total LSP protected : 0, Bandwidth in use : 0
Ingress (Primary)
4.4.4.4

```



```

From: 2.2.2.2, LSPstate: Up, LSPname: lsp3-Primary
Ingress FSM state: Operational
Establishment Time: 0s 17ms
SRLG configured in primary path are: 11 23
Setup priority: 7, Hold priority: 0
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
IGP-Shortcut: Disabled, LSP metric: 2
LSP Protection: facility
Fast-Reroute bandwidth : 0
Protection type desired: Link
Fast-Reroute Hop limit: 255
Fast-Reroute Setup priority: 7, Hold priority: 0
Label in: -, Label out: 24321,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5004, LSP Id: 2208, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 22.1.1.5, xel
Path refresh: 30 seconds (RR enabled) (due in 26421 seconds)
Resv lifetime: 157 seconds (due in 141 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Enabled
Admin Groups: none
Configured Path: sec-path1 (in use)
Configured Explicit Route Detail :
  4.4.4.4/32 loose
Session Explicit Route Detail :
  22.1.1.5/32 strict
  23.1.1.4/32 strict
Record route:
LP = 1 -> PLR's Downstream link is protected      PU = 1 -> Protection is in use on PLR
NP = 1 -> PLR's Downstream neighbor is protected  BP = 1 -> BW protection available at PLR
-----
IP Address      Label      (LP, PU, NP, BP)
-----
<self>
22.1.1.5        24321      ( 0,  0,  0,  0)
23.1.1.4        0           ( 0,  0,  0,  0)
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls

```

CSPF LSP:

```

RTR2#show cspf lsp
Lsp Id       : 0x13890899
OSPF ID      : 1
Ingress      : 2.2.2.2
Egress       : 4.4.4.4
Ext Tunnel ID : 2.2.2.2
LSP Type     : 0
Client ID    : 1
State        : 3
Setup Priority : 7
Hold Priority : 0
Hop Limit    : 255
Include Mask  : 0x0
Exclude Mask  : 0x0
Include All Mask: 0x0
LSP Metric   : 201
Path Constraint :
  11.1.1.3    strict

```

```

12.1.1.4      strict
Computed ERO  :
11.1.1.3
12.1.1.4

SRLG values of LSP :
11  12
Lsp Id       : 0x1389089c
OSPF ID      : 1
Ingress      : 2.2.2.2
Egress       : 4.4.4.4
Ext Tunnel ID : 2.2.2.2
LSP Type     : 0
Client ID    : 1
State        : 3
Setup Priority : 7
Hold Priority : 0
Hop Limit    : 255
Include Mask  : 0x0
Exclude Mask  : 0x0
Include All Mask: 0x0
LSP Metric   : 101
Path Constraint :
    4.4.4.4    loose
Exclude SRLG Constraint :
11  12
Computed ERO  :
    30.1.1.6
    31.1.1.4

SRLG values of LSP :
30  31
Lsp Id       : 0x138a089e
OSPF ID      : 1
Ingress      : 2.2.2.2
Egress       : 4.4.4.4
Ext Tunnel ID : 2.2.2.2
LSP Type     : 0
Client ID    : 1
State        : 3
Setup Priority : 7
Hold Priority : 0
Hop Limit    : 255
Include Mask  : 0x0
Exclude Mask  : 0x0
Include All Mask: 0x0
LSP Metric   : 201
Path Constraint :
    11.1.1.3    strict
    12.1.1.4    strict
Computed ERO  :
    11.1.1.3
    12.1.1.4

SRLG values of LSP :
11  12
Lsp Id       : 0x138a089e
OSPF ID      : 1
Ingress      : 30.1.1.2
Egress       : 4.4.4.4
Ext Tunnel ID : 2.2.2.2
LSP Type     : 0
Client ID    : 1
State        : 3
Setup Priority : 7
Hold Priority : 0
Hop Limit    : 255
Include Mask  : 0x0

```

```
Exclude Mask      : 0x0
Include All Mask: 0x0
LSP Metric        : 101
Exclude Path Constraint :
    11.1.1.3      exclude link
Exclude SRLG Constraint :
11
Computed ERO      :
    30.1.1.6
    31.1.1.4
```

```
SRLG values of LSP :
30  31
Lsp Id          : 0x138b089f
OSPF ID         : 1
Ingress         : 2.2.2.2
Egress          : 4.4.4.4
Ext Tunnel ID   : 2.2.2.2
LSP Type        : 2
Client ID       : 1
State           : 3
Setup Priority   : 7
Hold Priority    : 0
Hop Limit       : 255
Include Mask     : 0x0
Exclude Mask     : 0x0
Include All Mask: 0x0
LSP Metric      : 201
Path Constraint :
    11.1.1.3      strict
    12.1.1.4      strict
Computed ERO     :
    11.1.1.3
    12.1.1.4
```

```
SRLG values of LSP :
11  12
Lsp Id          : 0x138c08a0
OSPF ID         : 1
Ingress         : 2.2.2.2
Egress          : 4.4.4.4
Ext Tunnel ID   : 2.2.2.2
LSP Type        : 0
Client ID       : 1
State           : 3
Setup Priority   : 7
Hold Priority    : 0
Hop Limit       : 255
Include Mask     : 0x0
Exclude Mask     : 0x0
Include All Mask: 0x0
LSP Metric      : 2
Path Constraint :
    4.4.4.4      loose
Computed ERO     :
    22.1.1.5
    23.1.1.4
```

```
SRLG values of LSP :
11  23
```

RTR3 Configuration (Transit)

OSPF:

```
RTR3#show ip ospf neighbor
```

Total number of full neighbors: 3

OSPF process 1 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 2.2.2.2 | 1 | Full/Backup | 00:00:32 | 11.1.1.2 | xe7 | 0 |
| 4.4.4.4 | 1 | Full/Backup | 00:00:32 | 12.1.1.4 | ce52 | 0 |
| 5.5.5.5 | 1 | Full/Backup | 00:00:36 | 21.1.1.5 | xe5 | 0 |

RSVP:

```
RTR3#show rsvp session
```

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass

State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary

* indicates the session is active with local repair at one or more nodes

(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Type | LSPName | State |
|---|----------|------|-------------|----------------------|
| Uptime Rt Style Labelin Labelout DSType | | | | |
| 4.4.4.4 | 21.1.1.3 | DTR | lsp2-Detour | UP 01:12:50 1 1 SE - |
| 24320 | ELSP_CON | | | |

Total 1 displayed, Up 1, Down 0.

Transit RSVP:

| To | From | Type | LSPName | State |
|---|---------|----------|----------------|-----------------|
| Uptime Rt Style Labelin Labelout DSType | | | | |
| 4.4.4.4 | 2.2.2.2 | PRI | lsp1-Primary | UP 01:23:54 1 1 |
| SE 24960 | 0 | ELSP_CON | | |
| 4.4.4.4 | 2.2.2.2 | PRI | lsp2-Primary | UP 01:12:50 1 1 |
| SE 24961 | 0 | ELSP_CON | | |
| 4.4.4.4 | 2.2.2.2 | PRI | bypass1-Bypass | UP 01:07:35 1 1 |
| SE 24962 | 0 | ELSP_CON | | |

Total 3 displayed, Up 3, Down 0.

```
RTR3#show rsvp session detail
```

Transit

4.4.4.4

From: 2.2.2.2, LSPstate: Up, LSPname: lsp1-Primary

Transit upstream state: Operational, downstream state: Operational

Setup priority: 7, Hold priority: 0

IGP-Shortcut: Disabled, LSP metric: 65

LSP Protection: None

Label in: 24960, Label out: 0,

Tspec rate: 0, Fspec rate: 0

Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 2.2.2.2

Bind value: 0, Oper state: NA, Alloc mode: NA

Downstream: 12.1.1.4, ce52 Upstream: 11.1.1.2, xe7

Path refresh: 30 seconds (RR enabled) (due in 24976 seconds)

Path lifetime: 157 seconds (due in 112 seconds)

Resv refresh: 30 seconds (RR enabled) (due in 32606 seconds)

Resv lifetime: 157 seconds (due in 143 seconds)

RRO re-use as ERO: Disabled

Label Recording: Disabled

Admin Groups: Received Explicit Route Detail :

11.1.1.3/32 strict

12.1.1.4/32 strict

Session Explicit Route Detail :

12.1.1.4/32 strict

Record route:

```

-----
IP Address      Label
-----
11.1.1.2
<self>
12.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Transit
4.4.4.4
  From: 2.2.2.2, LSPstate: Up, LSPname: lsp2-Primary
  Transit upstream state: Operational, downstream state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: one-to-one
  Fast-Reroute bandwidth : 0
  Protection type desired: Link
  Fast-Reroute Hop limit: 255
  Fast-Reroute Setup priority: 7, Hold priority: 0
  Label in:  24961, Label out:  0,
  Tspecc rate: 0, Fspecc rate: 0
  Tunnel Id: 5002, LSP Id: 2206, Ext-Tunnel Id: 2.2.2.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 12.1.1.4, ce52 Upstream: 11.1.1.2, xe7
  Path refresh: 30 seconds (RR enabled) (due in 25624 seconds)
  Path lifetime: 157 seconds (due in 112 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 24416 seconds)
  Resv lifetime: 157 seconds (due in 143 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups:  Received Explicit Route Detail :
    11.1.1.3/32 strict
    12.1.1.4/32 strict
  Session Explicit Route Detail :
    12.1.1.4/32 strict
  Record route:
-----
IP Address      Label
-----
11.1.1.2
<self>
12.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Ingress (Detour)
4.4.4.4
  From: 21.1.1.3, LSPstate: Up, LSPname: lsp2-Detour
  Ingress FSM state: Operational
  Establishment Time: 0s 8ms
  SRLG configured in backup path are: 21 23
  Setup priority: 7, Hold priority: 0
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: OSPF
  IGP-Shortcut: Disabled, LSP metric: 2

```

```

LSP Protection: None
Label in: -, Label out: 24320,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5002, LSP Id: 2206, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 21.1.1.5, xe5
Path refresh: 30 seconds (RR enabled) (due in 25636 seconds)
Resv lifetime: 157 seconds (due in 153 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Exclude path detail:
  Exclude Link: 11.1.1.2
  Exclude Link: 12.1.1.4
Configured Path: none
Session Explicit Route Detail :
  21.1.1.5/32 strict
  23.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
<self>
21.1.1.5
23.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS    DSCP_value    EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Transit
4.4.4.4
From: 2.2.2.2, LSPstate: Up, LSPname: bypass1-Bypass
Transit upstream state: Operational, downstream state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in: 24962, Label out: 0,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5003, LSP Id: 2207, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 12.1.1.4, ce52 Upstream: 11.1.1.2, xe7
Path refresh: 30 seconds (RR enabled) (due in 25940 seconds)
Path lifetime: 157 seconds (due in 112 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 27142 seconds)
Resv lifetime: 157 seconds (due in 143 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: Received Explicit Route Detail :
  11.1.1.3/32 strict
  12.1.1.4/32 strict
Session Explicit Route Detail :
  12.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
11.1.1.2
<self>
12.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load

```

```

Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls

```

RTR4 Configuration (Egress)

OSPF:

```
RTR4#show ip ospf neighbor
```

```
Total number of full neighbors: 3
```

```
OSPF process 1 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 3.3.3.3 | 1 | Full/DR | 00:00:27 | 12.1.1.3 | ce52 | 0 |
| 5.5.5.5 | 1 | Full/Backup | 00:00:27 | 23.1.1.5 | xe9 | 0 |
| 6.6.6.6 | 1 | Full/Backup | 00:00:34 | 31.1.1.6 | ce50 | 0 |

RSVP:

```
RTR4#show rsvp session
```

```
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
```

```
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
```

```
* indicates the session is active with local repair at one or more nodes
```

```
(P) indicates the secondary-priority session is acting as primary
```

```
Egress RSVP:
```

| To | From | Type | LSPName | State |
|-----------|------------------------|---------------------------|---------|--------------|
| Uptime Rt | Style Labelin Labelout | DSType | | |
| 4.4.4.4 | 2.2.2.2 PRI | lsp1-Primary | UP | 01:28:26 1 1 |
| SE 0 | - ELSP_CON | | | |
| 4.4.4.4 | 2.2.2.2 PRI | lsp1-Secondary-Priority-1 | UP | 01:28:26 1 1 |
| SE 0 | - ELSP_CON | | | |
| 4.4.4.4 | 2.2.2.2 PRI | lsp2-Primary | UP | 01:17:22 1 1 |
| SE 0 | - ELSP_CON | | | |
| 4.4.4.4 | 30.1.1.2 PRI | lsp2-Detour | UP | 01:17:22 1 1 |
| SE 0 | - ELSP_CON | | | |
| 4.4.4.4 | 21.1.1.3 PRI | lsp2-Detour | UP | 01:17:22 1 1 |
| SE 0 | - ELSP_CON | | | |
| 4.4.4.4 | 2.2.2.2 PRI | bypass1-Bypass | UP | 01:12:07 1 1 |
| SE 0 | - ELSP_CON | | | |
| 4.4.4.4 | 2.2.2.2 PRI | lsp3-Primary | UP | 01:11:31 1 1 |
| SE 0 | - ELSP_CON | | | |

Total 7 displayed, Up 7, Down 0.

```
RTR4#show rsvp session detail
```

```
Egress
```

```
4.4.4.4
```

```
From: 2.2.2.2, LSPstate: Up, LSPname: lsp1-Primary
```

```
Egress FSM state: Operational
```

```
Setup priority: 7, Hold priority: 0
```

```
IGP-Shortcut: Disabled, LSP metric: 65
```

```
LSP Protection: None
```

```
Label in: 0, Label out: -,
```

```
Tspec rate: 0, Fspec rate: 0
```

```
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 2.2.2.2
```

```
Bind value: 0, Oper state: NA, Alloc mode: NA
```

```
Upstream: 12.1.1.3, ce52
```

```

Path lifetime: 157 seconds (due in 141 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 39085 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups:   Received Explicit Route Detail :
  12.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
11.1.1.2
12.1.1.3
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:   ELSP_CONFIG
CLASS      DSCP_value   EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Egress
4.4.4.4
From: 2.2.2.2, LSPstate: Up, LSPname: lsp1-Secondary-Priority-1
Egress FSM state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in:      0, Label out: -,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5001, LSP Id: 2204, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Upstream: 31.1.1.6, ce50
Path lifetime: 157 seconds (due in 141 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 9957 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups:   Received Explicit Route Detail :
  31.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
30.1.1.2
31.1.1.6
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:   ELSP_CONFIG
CLASS      DSCP_value   EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Egress
4.4.4.4
From: 2.2.2.2, LSPstate: Up, LSPname: lsp2-Primary
Egress FSM state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: one-to-one
Fast-Reroute bandwidth : 0
Protection type desired: Link
Fast-Reroute Hop limit: 255
Fast-Reroute Setup priority: 7, Hold priority: 0
Label in:      0, Label out: -,

```



```

Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5002, LSP Id: 2206, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Upstream: 12.1.1.3, ce52
Path lifetime: 157 seconds (due in 141 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 33578 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups:   Received Explicit Route Detail :
  12.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
11.1.1.2
12.1.1.3
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Egress
4.4.4.4
  From: 30.1.1.2, LSPstate: Up, LSPname: lsp2-Detour
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in:      0, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5002, LSP Id: 2206, Ext-Tunnel Id: 2.2.2.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Upstream: 31.1.1.6, ce50
  Path lifetime: 157 seconds (due in 141 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 27672 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups:   Received Explicit Route Detail :
    31.1.1.4/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  30.1.1.2
  31.1.1.6
  <self>
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type:  ELSP_CONFIG
  CLASS      DSCP_value      EXP_value
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Egress
4.4.4.4
  From: 21.1.1.3, LSPstate: Up, LSPname: lsp2-Detour
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in:      0, Label out: -,

```

```

Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5002, LSP Id: 2206, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Upstream: 23.1.1.5, xe9
Path lifetime: 157 seconds (due in 133 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 17352 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups:   Received Explicit Route Detail :
  23.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
21.1.1.3
23.1.1.5
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Egress
4.4.4.4
  From: 2.2.2.2, LSPstate: Up, LSPname: bypass1-Bypass
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in:      0, Label out: -,
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5003, LSP Id: 2207, Ext-Tunnel Id: 2.2.2.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Upstream: 12.1.1.3, ce52
  Path lifetime: 157 seconds (due in 141 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 22891 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups:   Received Explicit Route Detail :
    12.1.1.4/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  11.1.1.2
  12.1.1.3
  <self>
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type:  ELSP_CONFIG
  CLASS      DSCP_value      EXP_value
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Egress
4.4.4.4
  From: 2.2.2.2, LSPstate: Up, LSPname: lsp3-Primary
  Egress FSM state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: facility
  Fast-Reroute bandwidth : 0

```

```

Protection type desired: Link
Fast-Reroute Hop limit: 255
Fast-Reroute Setup priority: 7, Hold priority: 0
Label in:      0, Label out: -,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5004, LSP Id: 2208, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Upstream: 23.1.1.5, xe9
Path lifetime: 157 seconds (due in 133 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 33344 seconds)
RRO re-use as ERO: Disabled
Label Recording: Enabled
Admin Groups: Received Explicit Route Detail :
  23.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
22.1.1.2        24321
23.1.1.5        0
<self>
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type: ELSP_CONFIG
CLASS    DSCP_value    EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls

```

RTR5 Configuration

OSPF:

```
RTR5#show ip ospf neighbor
```

Total number of full neighbors: 3

OSPF process 1 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 3.3.3.3 | 1 | Full/DR | 00:00:34 | 21.1.1.3 | xe5 | 0 |
| 2.2.2.2 | 1 | Full/Backup | 00:00:30 | 22.1.1.2 | xe1 | 0 |
| 4.4.4.4 | 1 | Full/DR | 00:00:39 | 23.1.1.4 | xe9 | 0 |

RSVP:

```
RTR5#show rsvp session
```

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass

State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary

* indicates the session is active with local repair at one or more nodes

(P) indicates the secondary-priority session is acting as primary

Transit RSVP:

| To | From | Type | LSPName | State | |
|---------|----------|-------|--------------|----------|--------------|
| Uptime | Rt | Style | Labelin | Labelout | DSType |
| 4.4.4.4 | 21.1.1.3 | PRI | lsp2-Detour | UP | 01:23:45 1 1 |
| SE | 24320 | 0 | ELSP_CON | | |
| 4.4.4.4 | 2.2.2.2 | PRI | lsp3-Primary | UP | 01:17:54 1 1 |
| SE | 24321 | 0 | ELSP_CON | | |

Total 2 displayed, Up 2, Down 0.

```

RTR5#show rsvp session detail
Transit
4.4.4.4
  From: 21.1.1.3, LSPstate: Up, LSPname: lsp2-Detour
  Transit upstream state: Operational, downstream state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: 24320, Label out: 0,
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5002, LSP Id: 2206, Ext-Tunnel Id: 2.2.2.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 23.1.1.4, xe9 Upstream: 21.1.1.3, xe5
  Path refresh: 30 seconds (RR enabled) (due in 24907 seconds)
  Path lifetime: 157 seconds (due in 126 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 36355 seconds)
  Resv lifetime: 157 seconds (due in 144 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: Received Explicit Route Detail :
    21.1.1.5/32 strict
    23.1.1.4/32 strict
  Session Explicit Route Detail :
    23.1.1.4/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  21.1.1.3
  <self>
  23.1.1.4
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  LSP Type: ELSP_CONFIG
  CLASS DSCP_value EXP_value
  Last Recorded Error Code: None
  Last Recorded Error Value: None
  Node where Last Recorded Error originated: None
  Trunk Type: mpls
Transit
4.4.4.4
  From: 2.2.2.2, LSPstate: Up, LSPname: lsp3-Primary
  Transit upstream state: Operational, downstream state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: facility
  Fast-Reroute bandwidth : 0
  Protection type desired: Link
  Fast-Reroute Hop limit: 255
  Fast-Reroute Setup priority: 7, Hold priority: 0
  Label in: 24321, Label out: 0,
  Tspec rate: 0, Fspec rate: 0
  Tunnel Id: 5004, LSP Id: 2208, Ext-Tunnel Id: 2.2.2.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 23.1.1.4, xe9 Upstream: 22.1.1.2, xe1
  Path refresh: 30 seconds (RR enabled) (due in 25271 seconds)
  Path lifetime: 157 seconds (due in 148 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 38338 seconds)
  Resv lifetime: 157 seconds (due in 144 seconds)
  RRO re-use as ERO: Disabled
  Label Recording: Enabled
  Admin Groups: Received Explicit Route Detail :
    22.1.1.5/32 strict
    23.1.1.4/32 strict
  Session Explicit Route Detail :
    23.1.1.4/32 strict
  Record route:

```

```

-----
IP Address      Label
-----
22.1.1.2       24321
<self>
23.1.1.4       0
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls

```

RTR6 Configuration

OSPF:

```

RTR6#show ip ospf neighbor
Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface      Instance ID
2.2.2.2        1    Full/Backup     00:00:32    30.1.1.2       xe5            0
4.4.4.4        1    Full/DR         00:00:38    31.1.1.4       ce50           0

```

RSVP:

```

RTR6#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

```

Transit RSVP:
To      From      Type      LSPName      State
Uptime  Rt  Style  Labelin  Labelout  DStype
4.4.4.4  2.2.2.2  PRI      lsp1-Secondary-Priority-1  UP    01:38:16  1 1
SE      24320  0        ELSP_CON
4.4.4.4  30.1.1.2  PRI      lsp2-Detour   UP    01:27:12  1 1
SE      24321  0        ELSP_CON
Total 2 displayed, Up 2, Down 0.

```

```

RTR6#show rsvp session detail
Transit
4.4.4.4
  From: 2.2.2.2, LSPstate: Up, LSPname: lsp1-Secondary-Priority-1
  Transit upstream state: Operational, downstream state: Operational
  Setup priority: 7, Hold priority: 0
  IGP-Shortcut: Disabled, LSP metric: 65
  LSP Protection: None
  Label in: 24320, Label out: 0,
  Tspect rate: 0, Fspect rate: 0
  Tunnel Id: 5001, LSP Id: 2204, Ext-Tunnel Id: 2.2.2.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 31.1.1.4, ce50 Upstream: 30.1.1.2, xe5
  Path refresh: 30 seconds (RR enabled) (due in 24080 seconds)
  Path lifetime: 157 seconds (due in 149 seconds)
  Resv refresh: 30 seconds (RR enabled) (due in 37411 seconds)
  Resv lifetime: 157 seconds (due in 150 seconds)
  RRO re-use as ERO: Disabled

```

```

Label Recording: Disabled
Admin Groups:   Received Explicit Route Detail :
  30.1.1.6/32 strict
  31.1.1.4/32 strict
Session Explicit Route Detail :
  31.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
30.1.1.2
<self>
31.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls
Transit
4.4.4.4
From: 30.1.1.2, LSPstate: Up, LSPname: lsp2-Detour
Transit upstream state: Operational, downstream state: Operational
Setup priority: 7, Hold priority: 0
IGP-Shortcut: Disabled, LSP metric: 65
LSP Protection: None
Label in:      24321, Label out:      0,
Tspec rate: 0, Fspec rate: 0
Tunnel Id: 5002, LSP Id: 2206, Ext-Tunnel Id: 2.2.2.2
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 31.1.1.4, ce50 Upstream: 30.1.1.2, xe5
Path refresh: 30 seconds (RR enabled) (due in 24724 seconds)
Path lifetime: 157 seconds (due in 149 seconds)
Resv refresh: 30 seconds (RR enabled) (due in 32187 seconds)
Resv lifetime: 157 seconds (due in 150 seconds)
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups:   Received Explicit Route Detail :
  30.1.1.6/32 strict
  31.1.1.4/32 strict
Session Explicit Route Detail :
  31.1.1.4/32 strict
Record route:
-----
IP Address      Label
-----
30.1.1.2
<self>
31.1.1.4
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
LSP Type:  ELSP_CONFIG
CLASS      DSCP_value      EXP_value
Last Recorded Error Code: None
Last Recorded Error Value: None
Node where Last Recorded Error originated: None
Trunk Type: mpls

```

RSVP SRLG with ISIS IGP

Overview

Shared Risk Link Groups (SRLG) allow the network to compute backup secondary LSPs or FRR (fast-reroute) LSPs that are disjoint from the primary LSP. Links grouped under the same SRLG share a common underlying risk—for example, fibers routed through the same conduit or multiple wavelengths carried on the same fiber. Because a failure on one link in an SRLG is likely to affect all links in that group, SRLG-aware path computation ensures that alternate LSPs avoid shared risks and maintain path diversity.

Characteristics

SRLG groups identify which links share the same risk. As an Interior Gateway Protocol (IGP), IS-IS distributes network topology information—including link attributes—throughout the network. To advertise SRLG membership, IS-IS uses the Shared Risk Link Group TLV (Type 138).

CSPF (Constrained Shortest Path First) computes LSP paths based on multiple constraints, taking into account not only the network topology but also link and LSP attributes. The CSPF engine uses information from both the signaling and routing modules to determine an optimal path while efficiently balancing network load.

SRLG for Secondary Sessions

- The SRLG of an LSP path is the union of SRLGs across all its links.
- When computing a secondary LSP, CSPF attempts to select a path that is SRLG-disjoint from the primary path.
- This avoids correlated failures that could take down both primary and secondary LSPs.

SRLG for Fast Reroute (FRR)

RSVP supports two FRR modes:

- One-to-One Protection (Detours)
- Facility Backup Protection (Bypass Tunnels)

SRLG constraints prevent detours or bypass tunnels from using links with the same SRLG as the protected link.

Detour LSPs

- Each primary LSP has its own detour.
- Detour path selection avoids links sharing SRLGs with the protected link.

Auto Bypass

- Dynamically created bypass tunnels protect multiple LSPs.
- Auto-bypass path computation avoids links that share SRLGs with protected interfaces.

Manual Bypass

- Manually configured bypass tunnels can be mapped to primary sessions.
- Mapping requires SRLG exclusivity between the bypass tunnel and the protected link.

Benefits

- Ensures secondary or FRR LSPs do not share risk with primary paths.
 - Reduces probability that a single failure impacts both primary and backup LSPs.
 - CSPF automatically avoids links with common risks.
 - Enhances primary/secondary and FRR path diversity.
-

SRLG with ISIS IGPC Configuration

Prerequisites

- Loopback IPs configured on all nodes (used as LSP endpoints).
- IS-IS enabled and fully operational across the entire network.
- RSVP-TE trunks established between Head and Tail nodes.

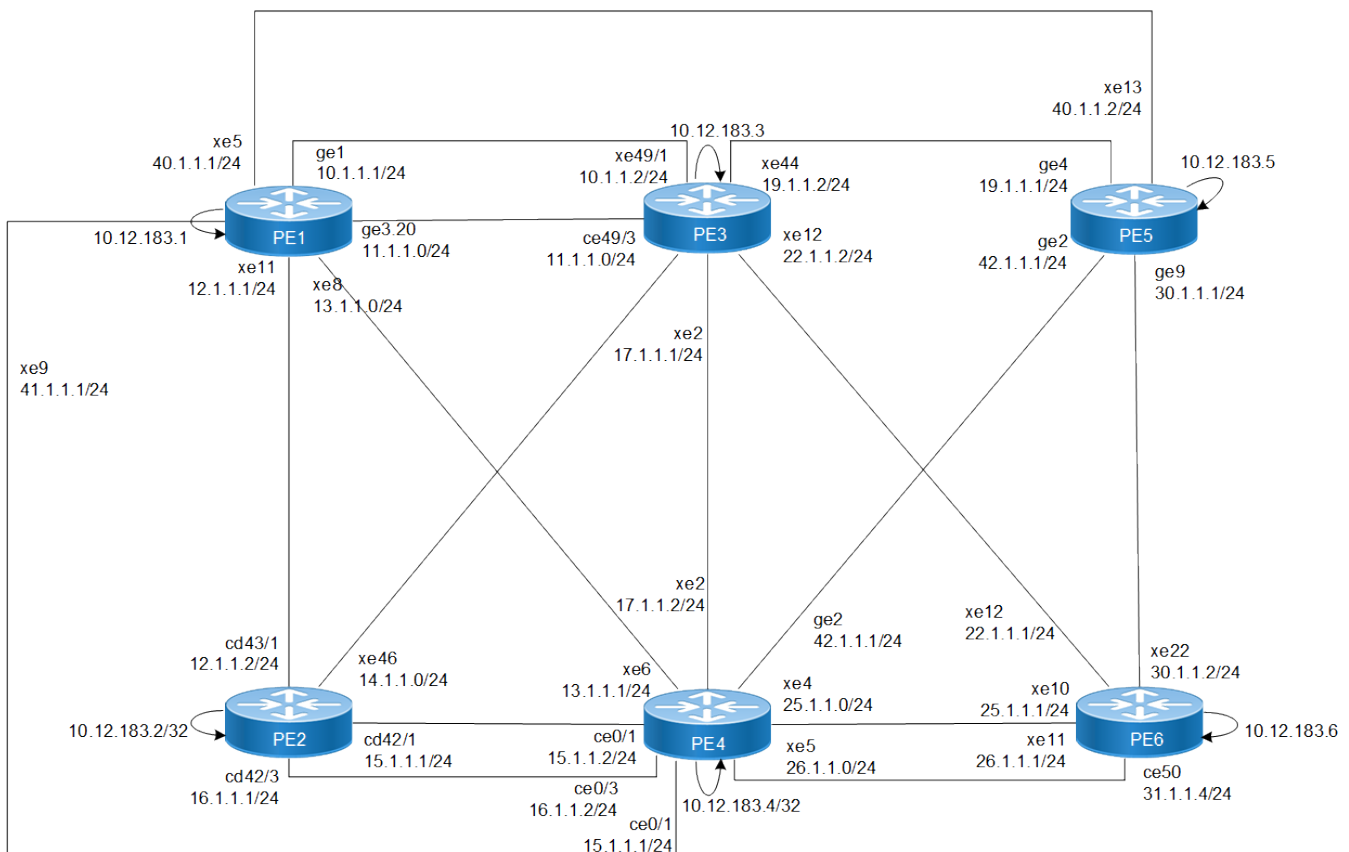
Topology

The figure below illustrates six Provider Edge (PE) routers—PE1 through PE6—interconnected in a partially-meshed MPLS/IP core. Each router also has a loopback interface for routing, and multiple point-to-point links interconnect the nodes in a structured pattern designed.

Node Summary

Each PE router has:

- A primary loopback address (lo: 10.12.183.x)
- Multiple routed point-to-point interfaces (xe, ge, ce, cd)
- Access to one or more backbone subnets and inter-PE adjacency links



Configuration Steps

The configurations enable:

- RSVP-TE on all PE-to-PE links
- ISIS L1 as the Interior Gateway Protocol
- Traffic Engineering (TE) extensions
- Loopback router-IDs for MPLS/IGP
- Optional PIM/BFD configurations on specific interfaces



Note: Before configuration meet all [Prerequisites \(page 1148\)](#).

PE1

1. Enable RSVP on the interface ge1.

```
(config)#interface ge1
(config-if)#load-interval 30
(config-if)#ip address 10.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

2. Enable RSVP on the interface ge3.

```
(config)#interface ge3.20
(config-if)#encapsulation dot1q 20
(config-if)#load-interval 30
(config-if)#ip address 11.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

3. Configure loopback interface.

```
(config)#interface lo
(config-if)#ip address 10.12.183.1/32 secondary
(config-if)#ip router isis ISIS-IGP-100
(config-if)#commit
(config-if)#exit
```

4. Enable RSVP on the interface xe5.

```
(config)#interface xe5
(config-if)#load-interval 30
(config-if)#ip address 40.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

5. Enable RSVP on the interface xe8.

```
(config)#interface xe8
(config-if)#load-interval 30
(config-if)#ip address 13.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

6. Enable RSVP on the interface xe9.

```
(config)#interface xe9
(config-if)#load-interval 30
(config-if)#ip address 41.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

7. Enable RSVP on the interface xe11.

```
(config)#interface xe11
(config-if)#load-interval 30
(config-if)#ip address 12.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

8. Enable IS-IS as the IGP.

```
(config)#router isis ISIS-IGP-100
(config-router)#is-type level-1
(config-router)#ignore-lsp-errors
(config-router)#lsp-gen-interval 5
(config-router)#spf-interval-exp level-1 50 2000
(config-router)#metric-style wide
(config-router)#mpls traffic-eng router-id 10.12.183.1
(config-router)#mpls traffic-eng level-1
(config-router)#capability cspf
(config-router)#dynamic-hostname
(config-router)#bfd all-interfaces
(config-router)#net 49.0001.0000.0000.0001.00
(config-if)#commit
(config-if)#exit
```

PE2

1. Enable RSVP on the interface cd42/1.

```
(config)#interface cd42/1
(config-if)#load-interval 30
(config-if)#ip address 15.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

2. Enable RSVP on the interface cd42/3.

```
(config)#interface cd42/3
(config-if)#load-interval 30
(config-if)#ip address 16.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

3. Enable RSVP on the interface cd43/1.

```
(config)#interface cd43/1
(config-if)#load-interval 30
(config-if)#ip address 12.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

4. Enable RSVP on the interface ce18.

```
(config)#interface ce18
(config-if)#load-interval 30
(config-if)#ip address 33.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

5. Configure loopback interface.

```
(config)#interface lo
```

```
(config-if)#ip address 10.12.183.2/32 secondary
(config-if)#ip router isis ISIS-IGP-100
(config-if)#commit
(config-if)#exit
```

6. Enable RSVP on the interface xe46.

```
(config)#interface xe46
(config-if)#load-interval 30
(config-if)#ip address 14.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

7. Enable IS-IS as the IGP.

```
(config)#router isis ISIS-IGP-100
(config-router)#is-type level-1
(config-router)#ignore-lsp-errors
(config-router)#lsp-gen-interval 5
(config-router)#spf-interval-exp level-1 50 2000
(config-router)#metric-style wide
(config-router)#mpls traffic-eng router-id 10.12.183.2
(config-router)#mpls traffic-eng level-1
(config-router)#capability cspf
(config-router)#dynamic-hostname
(config-router)#bfd all-interfaces
(config-router)#net 49.0001.0000.0000.0002.00
(config-if)#commit
(config-if)#exit
```

PE3

1. Enable RSVP on the interface ce49/1.

```
(config)#interface ce49/1
(config-if)#speed 1g
(config-if)#load-interval 30
(config-if)#ip address 10.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

2. Enable RSVP on the interface ce49/3.

```
(config)#interface ce49/3.20
(config-if)#encapsulation dot1q 20
(config-if)#load-interval 30
(config-if)#ip address 11.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

3. Configure loopback interface.

```
(config)#interface lo
(config-if)#ip address 10.12.183.3/32 secondary
(config-if)#ip router isis ISIS-IGP-100
(config-if)#commit
(config-if)#exit
```

4. Configure the interface.

```
(config)#interface xe1
(config-if)#load-interval 30
(config-if)#ip address 21.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#ip pim sparse-mode
(config-if)#bfd interval 3 minrx 3 multiplier 3
(config-if)#commit
(config-if)#exit
```

5. Enable RSVP on the interface xe2.

```
(config)#interface xe2
(config-if)#load-interval 30
(config-if)#ip address 17.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

6. Enable RSVP on the interface xe3.

```
(config)#interface xe3
(config-if)#load-interval 30
(config-if)#ip address 14.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

7. Enable RSVP on the interface xe12.

```
(config)#interface xe12
(config-if)#load-interval 30
(config-if)#ip address 22.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

8. Enable RSVP on the interface xe44.

```
(config)#interface xe44
(config-if)#load-interval 30
(config-if)#ip address 19.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

9. Enable IS-IS as the IGP.

```
(config)#router isis ISIS-IGP-100
(config-router)#is-type level-1
(config-router)#ignore-lsp-errors
(config-router)#lsp-gen-interval 5
(config-router)#spf-interval-exp level-1 50 2000
(config-router)#metric-style wide
```

```
(config-router)#mpls traffic-eng router-id 10.12.183.3
(config-router)#mpls traffic-eng level-1
(config-router)#capability cspf
(config-router)#dynamic-hostname
(config-router)#bfd all-interfaces
(config-router)#net 49.0001.0000.0000.0003.00
(config-if)#commit
(config-if)#exit
```

PE4

1. Enable RSVP on the interface ce0/1.

```
(config)#interface ce0/1
(config-if)#load-interval 30
(config-if)#ip address 15.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

2. Enable RSVP on the interface ce0/3.

```
(config)#interface ce0/3
(config-if)#load-interval 30
(config-if)#ip address 16.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

3. Enable RSVP on the interface ge1.

```
(config)#interface ge1
(config-if)#load-interval 30
(config-if)#ip address 24.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

4. Enable RSVP on the interface ge2.

```
(config)#interface ge2
(config-if)#load-interval 30
(config-if)#ip address 42.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

5. Configure loopback interface.

```
(config)#interface lo
(config-if)#ip address 10.12.183.4/32 secondary
(config-if)#ip router isis ISIS-IGP-100
(config-if)#commit
(config-if)#exit
```

6. Configure the interface.

```
(config)#interface xe0
(config-if)#load-interval 30
(config-if)#ip address 41.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

7. Enable RSVP on the interface xe2.

```
(config)#interface xe2
(config-if)#load-interval 30
(config-if)#ip address 17.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

8. Enable RSVP on the interface xe4.

```
(config)#interface xe4
(config-if)#load-interval 30
(config-if)#ip address 25.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

9. Enable RSVP on the interface xe5.

```
(config)#interface xe5
(config-if)#load-interval 30
(config-if)#ip address 26.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

10 Enable RSVP on the interface xe6.

```
(config)#interface xe6
(config-if)#load-interval 30
(config-if)#ip address 13.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

11 Enable IS-IS as the IGP.

```
(config)#router isis ISIS-IGP-100
(config-router)#is-type level-1
(config-router)#ignore-lsp-errors
(config-router)#lsp-gen-interval 5
(config-router)#spf-interval-exp level-1 50 2000
(config-router)#metric-style wide
(config-router)#mpls traffic-eng router-id 10.12.183.4
(config-router)#mpls traffic-eng level-1
(config-router)#capability cspf
(config-router)#dynamic-hostname
```

```
(config-router)#bfd all-interfaces
(config-router)#net 49.0001.0000.0000.0004.00
(config-if)#commit
(config-if)#exit
```

PE5

1. Enable RSVP on the interface ge2.

```
(config)#interface ge2
(config-if)#load-interval 30
(config-if)#ip address 42.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

2. Enable RSVP on the interface ge4.

```
(config)#interface ge4
(config-if)#load-interval 30
(config-if)#ip address 19.1.1.1/24
(config-if)#label-switching
(config-if)#holddown-timer 900
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

3. Enable RSVP on the interface ge7.

```
(config)#interface ge7.20
(config)#encapsulation dot1q 20
(config-if)#load-interval 30
(config-if)#ip address 20.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

4. Enable RSVP on the interface ge8.

```
(config)#interface ge8
(config-if)#load-interval 30
(config-if)#ip address 29.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

5. Enable RSVP on the interface ge9.

```
(config)#interface ge9
(config-if)#ip address 30.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

6. Configure loopback interface.


```
(config)#interface lo
(config-if)#ip address 10.12.183.5/32 secondary
(config-if)#ip router isis ISIS-IGP-100
(config-if)#commit
(config-if)#exit
```

7. Enable RSVP on the interface xe13.

```
(config)#interface xe13
(config-if)#load-interval 30
(config-if)#ip address 40.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

8. Enable IS-IS as the IGP.

```
(config)#router isis ISIS-IGP-100
(config-router)#is-type level-1
(config-router)#ignore-lsp-errors
(config-router)#lsp-gen-interval 5
(config-router)#spf-interval-exp level-1 50 2000
(config-router)#metric-style wide
(config-router)#mpls traffic-eng router-id 10.12.183.5
(config-router)#mpls traffic-eng level-1
(config-router)#capability cspf
(config-router)#dynamic-hostname
(config-router)#bfd all-interfaces
(config-router)#net 49.0001.0000.0000.0005.00
(config-if)#commit
(config-if)#exit
```

PE6

1. Configure loopback interface.

```
(config)#interface lo
(config-if)#ip address 10.12.183.6/32 secondary
(config-if)#ip router isis ISIS-IGP-100
(config-if)#commit
(config-if)#exit
```

2. Enable RSVP on the interface xe10.

```
(config)#interface xe10
(config-if)#load-interval 30
(config-if)#ip address 25.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

3. Enable RSVP on the interface xe11.

```
(config)#interface xe11
(config-if)#speed 10g
(config-if)#load-interval 30
(config-if)#ip address 26.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

4. Enable RSVP on the interface xe12.

```
(config)#interface xe12
(config-if)#load-interval 30
(config-if)#ip address 22.1.1.1/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

5. Enable RSVP on the interface xe22.

```
(config)#interface xe22
(config-if)#ip address 30.1.1.2/24
(config-if)#label-switching
(config-if)#isis network point-to-point
(config-if)#ip router isis ISIS-IGP-100
(config-if)#enable-rsvp
(config-if)#commit
(config-if)#exit
```

6. Enable IS-IS as the IGP.

```
(config)#router isis ISIS-IGP-100
(config-router)#is-type level-1
(config-router)#ignore-lsp-errors
(config-router)#lsp-gen-interval 5
(config-router)#spf-interval-exp level-1 50 2000
(config-router)#metric-style wide
(config-router)#mpls traffic-eng router-id 10.12.183.6
(config-router)#mpls traffic-eng level-1
(config-router)#capability cspf
(config-router)#dynamic-hostname
(config-router)#bfd all-interfaces
(config-router)#net 49.0001.0000.0000.0006.00
(config-if)#commit
(config-if)#exit
```

SRLG with ISIS IGP Running Configuration**PE1**

```
!
router rsvp
!
interface ge1
load-interval 30
ip address 10.1.1.1/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface ge3.20
encapsulation dot1q 20
load-interval 30
ip address 11.1.1.1/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface lo
ip address 10.12.183.1/32 secondary
ip router isis ISIS-IGP-100
```

```
!  
interface xe5  
  load-interval 30  
  ip address 40.1.1.1/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface xe8  
  load-interval 30  
  ip address 13.1.1.1/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface xe9  
  load-interval 30  
  ip address 41.1.1.1/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface xell  
  load-interval 30  
  ip address 12.1.1.1/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
  exit  
!  
!  
router isis ISIS-IGP-100  
  is-type level-1  
  ignore-lsp-errors  
  lsp-gen-interval 5  
  spf-interval-exp level-1 50 2000  
  metric-style wide  
  mpls traffic-eng router-id 10.12.183.1  
  mpls traffic-eng level-1  
  capability cspf  
  dynamic-hostname  
  bfd all-interfaces  
  net 49.0001.0000.0000.0001.00  
!  
  exit  
!  
end
```

PE2

```
!  
router rsvp  
!  
interface cd42/1  
  load-interval 30  
  ip address 15.1.1.1/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface cd42/3  
  load-interval 30
```

```
ip address 16.1.1.1/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface cd43/1
load-interval 30
ip address 12.1.1.2/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface ce18
load-interval 30
ip address 33.1.1.1/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface lo
ip address 10.12.183.2/32 secondary
ip router isis ISIS-IGP-100
!
interface xe46
load-interval 30
ip address 14.1.1.1/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
exit
!
router isis ISIS-IGP-100
is-type level-1
ignore-lsp-errors
lsp-gen-interval 5
spf-interval-exp level-1 50 2000
metric-style wide
mpls traffic-eng router-id 10.12.183.2
mpls traffic-eng level-1
capability cspf
dynamic-hostname
bfd all-interfaces
net 49.0001.0000.0000.0002.00
!
exit
end
```

PE3

```
!
router rsvp
!
interface ce49/1
speed 1g
load-interval 30
ip address 10.1.1.2/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface ce49/3.20
encapsulation dot1q 20
```

```
load-interval 30
ip address 11.1.1.2/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface lo
ip address 10.12.183.3/32 secondary
ip router isis ISIS-IGP-100
!
interface xel
load-interval 30
ip address 21.1.1.2/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
ip pim sparse-mode
bfd interval 3 minrx 3 multiplier 3
!
interface xe2
load-interval 30
ip address 17.1.1.1/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface xe3
load-interval 30
ip address 14.1.1.2/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface xel2
load-interval 30
ip address 22.1.1.2/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface xe44
speed lg
load-interval 30
ip address 19.1.1.2/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
exit
!!
router isis ISIS-IGP-100
is-type level-1
ignore-lsp-errors
lsp-gen-interval 5
spf-interval-exp level-1 50 2000
metric-style wide
mpls traffic-eng router-id 10.12.183.3
mpls traffic-eng level-1
capability cspf
dynamic-hostname
bfd all-interfaces
net 49.0001.0000.0000.0003.00
!
```

```
exit
!  
end
```

PE4

```
!  
router rsvp  
!  
interface ce0/1  
  load-interval 30  
  ip address 15.1.1.2/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface ce0/3  
  load-interval 30  
  ip address 16.1.1.2/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface gel  
  load-interval 30  
  ip address 24.1.1.2/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface ge2  
  load-interval 30  
  ip address 42.1.1.2/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface lo  
  ip address 10.12.183.4/32 secondary  
  ip router isis ISIS-IGP-100  
!  
interface xe0  
  load-interval 30  
  ip address 41.1.1.2/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface xe2  
  load-interval 30  
  ip address 17.1.1.2/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface xe4  
  load-interval 30  
  ip address 25.1.1.2/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp
```

```
!  
interface xe5  
  load-interval 30  
  ip address 26.1.1.2/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface xe6  
  load-interval 30  
  ip address 13.1.1.2/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
  exit  
!  
!  
router isis ISIS-IGP-100  
  is-type level-1  
  ignore-lsp-errors  
  lsp-gen-interval 5  
  spf-interval-exp level-1 50 2000  
  metric-style wide  
  mpls traffic-eng router-id 10.12.183.4  
  mpls traffic-eng level-1  
  capability cspf  
  dynamic-hostname  
  bfd all-interfaces  
  net 49.0001.0000.0000.0004.00  
!  
  exit  
!  
end
```

PE5

```
!  
router rsvp  
!  
interface ge2  
  load-interval 30  
  ip address 42.1.1.1/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface ge4  
  load-interval 30  
  ip address 19.1.1.1/24  
  label-switching  
  holddown-timer 900  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!  
interface ge7.20  
  encapsulation dot1q 20  
  load-interval 30  
  ip address 20.1.1.1/24  
  label-switching  
  isis network point-to-point  
  ip router isis ISIS-IGP-100  
  enable-rsvp  
!
```

```

interface ge8
  load-interval 30
  ip address 29.1.1.1/24
  label-switching
  isis network point-to-point
  ip router isis ISIS-IGP-100
  enable-rsvp
!
interface ge9
  ip address 30.1.1.1/24
  label-switching
  isis network point-to-point
  ip router isis ISIS-IGP-100
  enable-rsvp
!
interface lo
  ip address 10.12.183.5/32 secondary
  ip router isis ISIS-IGP-100
!
interface xe13
  load-interval 30
  ip address 40.1.1.2/24
  label-switching
  isis network point-to-point
  ip router isis ISIS-IGP-100
  enable-rsvp
!
  exit
!
router isis ISIS-IGP-100
  is-type level-1
  ignore-lsp-errors
  lsp-gen-interval 5
  spf-interval-exp level-1 50 2000
  metric-style wide
  mpls traffic-eng router-id 10.12.183.5
  mpls traffic-eng level-1
  capability cspf
  dynamic-hostname
  bfd all-interfaces
  net 49.0001.0000.0000.0005.00
!
  exit
!
end

```

PE6

```

!
router rsvp
!
interface lo
  ip address 10.12.183.6/32 secondary
  ip router isis ISIS-IGP-100
!
interface xe10
  load-interval 30
  ip address 25.1.1.1/24
  label-switching
  isis network point-to-point
  ip router isis ISIS-IGP-100
  enable-rsvp
!
interface xe11
  speed 10g
  load-interval 30
  ip address 26.1.1.1/24
  label-switching

```



```
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface xe12
load-interval 30
ip address 22.1.1.1/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
interface xe22
ip address 30.1.1.2/24
label-switching
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!
exit
!
!
router isis ISIS-IGP-100
is-type level-1
ignore-lsp-errors
lsp-gen-interval 5
spf-interval-exp level-1 50 2000
metric-style wide
mpls traffic-eng router-id 10.12.183.6
mpls traffic-eng level-1
capability cspf
dynamic-hostname
bfd all-interfaces
net 49.0001.0000.0000.0006.00
!
exit
!
end
```

Validation

PE

```
PE1#sh running-config isis
!
!
router isis ISIS-IGP-100
is-type level-1
ignore-lsp-errors
lsp-gen-interval 5
spf-interval-exp level-1 50 2000
metric-style wide
mpls traffic-eng router-id 10.12.183.1
mpls traffic-eng level-1
capability cspf
dynamic-hostname
bfd all-interfaces
net 49.0001.0000.0000.0001.00
!
interface ge1
isis network point-to-point
ip router isis ISIS-IGP-100
!
interface ge2
isis network point-to-point
ip router isis ISIS-IGP-100
```

```

!
interface lo
  ip router isis ISIS-IGP-100
!
interface xe5
  isis network point-to-point
  ip router isis ISIS-IGP-100
!
interface xe8
  isis network point-to-point
  ip router isis ISIS-IGP-100
!
interface xe9
  isis network point-to-point
  ip router isis ISIS-IGP-100
!
interface xe11
  isis network point-to-point
  ip router isis ISIS-IGP-100
!
interface ge3.20
  isis network point-to-point
  ip router isis ISIS-IGP-100
!
!
PE1#

```

```
PE1#sh ip route
```

```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

```

```
IP Route Table for VRF "default"
```

```

C          10.12.183.1/32 is directly connected, lo, installed 02:55:55, last update 02:55:55 ago
i L1       10.12.183.2/32 [115/20] via 12.1.1.2, xe11, installed 01:40:06, last update 01:40:06 ago
i L1       10.12.183.3/32 [115/20] via 11.1.1.2, ge3.20, installed 01:40:06, last update 01:36:40
ago
i L1       10.12.183.4/32 [115/20] via 41.1.1.2, xe9, installed 01:40:06, last update 01:40:01 ago
           [115/20] via 13.1.1.2, xe8
i L1       10.12.183.5/32 [115/20] via 40.1.1.2, xe5, installed 01:40:01, last update 01:40:01 ago
i L1       10.12.183.6/32 [115/30] via 11.1.1.2, ge3.20, installed 01:40:06, last update 01:36:40
ago
           [115/30] via 41.1.1.2, xe9
           [115/30] via 13.1.1.2, xe8
           [115/30] via 40.1.1.2, xe5
C          11.1.1.0/24 is directly connected, ge3.20, installed 01:37:21, last update 01:37:21 ago
C          12.1.1.0/24 is directly connected, xe11, installed 02:53:43, last update 02:53:43 ago
C          13.1.1.0/24 is directly connected, xe8, installed 02:51:56, last update 02:51:56 ago
i L1       14.1.1.0/24 [115/20] via 11.1.1.2, ge3.20, installed 01:02:28, last update 01:02:28 ago
           [115/20] via 13.1.1.2, xe11
i L1       15.1.1.0/24 [115/20] via 41.1.1.2, xe9, installed 01:40:06, last update 01:40:01 ago
           [115/20] via 13.1.1.2, xe8
           [115/20] via 12.1.1.2, xe11
i L1       16.1.1.0/24 [115/20] via 41.1.1.2, xe9, installed 01:40:06, last update 01:40:01 ago
           [115/20] via 13.1.1.2, xe8
           [115/20] via 12.1.1.2, xe11
i L1       17.1.1.0/24 [115/20] via 11.1.1.2, ge3.20, installed 01:40:06, last update 01:36:40 ago
           [115/20] via 41.1.1.2, xe9
           [115/20] via 13.1.1.2, xe8
i L1       19.1.1.0/24 [115/20] via 11.1.1.2, ge3.20, installed 01:40:01, last update 01:36:40 ago
           [115/20] via 40.1.1.2, xe5
i L1       22.1.1.0/24 [115/20] via 11.1.1.2, ge3.20, installed 01:40:06, last update 01:36:40 ago

```

```

i L1      25.1.1.0/24 [115/20] via 41.1.1.2, xe9, installed 01:40:06, last update 01:40:01 ago
           [115/20] via 13.1.1.2, xe8
i L1      26.1.1.0/24 [115/20] via 41.1.1.2, xe9, installed 01:40:06, last update 01:40:01 ago
           [115/20] via 13.1.1.2, xe8
i L1      30.1.1.0/24 [115/20] via 40.1.1.2, xe5, installed 01:40:01, last update 01:40:01 ago
C         40.1.1.0/24 is directly connected, xe5, installed 02:55:55, last update 02:55:55 ago
C         41.1.1.0/24 is directly connected, xe9, installed 01:40:12, last update 01:40:12 ago
i L1      42.1.1.0/24 [115/20] via 41.1.1.2, xe9, installed 01:40:06, last update 01:40:01 ago
           [115/20] via 13.1.1.2, xe8
           [115/20] via 40.1.1.2, xe5
C         127.0.0.0/8 is directly connected, lo, installed 03:35:28, last update 03:35:28 ago

Gateway of last resort is not set
PE1#

```

Detour(Frr)

```

PE1#

rsvp-path PE1-PE6-via-PE2-PE4 mpls
 12.1.1.2 strict
 15.1.1.2 strict
 25.1.1.1 strict
!
!
!
rsvp-trunk PE1-PE6 ipv4
 reoptimize
 primary fast-reroute protection one-to-one
 primary path PE1-PE6-via-PE2-PE4
 to 10.12.183.6
!

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime    Rt   Style  Labelin Labelout
10.12.183.6 10.12.183.1 5001     2201     PRI    PE1-PE6-         -
Primary          UP    00:07:33 1 1 SE    -          24320
10.12.183.6 40.1.1.1    5001     2201     DTR    PE1-PE6-         -
Detour          UP    00:07:33 1 1 SE    -          24320
Total 2 displayed, Up 2, Down 0.

PE1#

PE1#sh rsvp session PE1-PE6-Primary
Ingress (Primary)
10.12.183.6
  From: 10.12.183.1, LSPstate: Up, LSPname: PE1-PE6-Primary
  Ingress FSM state: Operational
  Establishment Time: 0s 33ms
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 30
  LDP Tunneling : Disabled
  LSP Protection: one-to-one
  Fast-Reroute bandwidth : 0
  Protection type desired: Link
  Fast-Reroute Hop limit: 255
  Fast-Reroute Setup priority: 7, Hold priority: 0

```

```

Label in: -, Label out: 24320,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 10.12.183.1
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 12.1.1.2, xe11
Path refresh: 30 seconds (RR enabled) (due in 29548 seconds)
Resv lifetime: 157 seconds (due in 138 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: PE1-PE6-via-PE2-PE4 (in use)
Configured Explicit Route Detail :
  12.1.1.2/32 strict
  15.1.1.2/32 strict
  25.1.1.1/32 strict
Session Explicit Route Detail :
  12.1.1.2/32 strict
  15.1.1.2/32 strict
  25.1.1.1/32 strict
Record route:
LP = 1 -> PLR's Downstream link is protected      PU = 1 -> Protection is in use on PLR
NP = 1 -> PLR's Downstream neighbor is protected  BP = 1 -> BW protection available at PLR
-----
IP Address      Label      (LP, PU, NP, BP)
-----
<self>
12.1.1.2        ( 1, 0, 0, 0)
15.1.1.2        ( 1, 0, 0, 0)
25.1.1.1        ( 0, 0, 0, 0)
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500

```

```

PE1#sh rsvp session PE1-PE6-Detour
Ingress (Detour)
10.12.183.6
  From: 40.1.1.1, LSPstate: Up, LSPname: PE1-PE6-Detour
  Ingress FSM state: Operational
  Establishment Time: 0s 9ms
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled (Implicit), Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LDP Tunneling : Disabled
  LSP Protection: None
  Label in: -, Label out: 24320,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 40.1.1.2, xe5
  Path refresh: 30 seconds (RR enabled) (due in 29534 seconds)
  Resv lifetime: 157 seconds (due in 155 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Exclude path detail:
    Exclude Link: 12.1.1.2
  Configured Path: none
  Session Explicit Route Detail :
    40.1.1.2/32 strict
    30.1.1.2/32 strict
  Record route:
  -----
  IP Address      Label

```

```

-----
<self>
40.1.1.2
30.1.1.2
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500

```

Configure the SRLG values on both primary and detour path

```

PE1(config)#int xe11
PE1(config-if)#mpls traffic-eng srlg 10
PE1(config-if)#commit
PE1(config-if)#int xe5
PE1(config-if)#mpls traffic-eng srlg 10
PE1(config-if)#commit
PE1(config-if)#end

```

Configure the srlg-disjoint as preferred

```

PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint preferred
PE1(config-router)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

```

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime    Rt   Style  Labelin  Labelout
10.12.183.6  10.12.183.1  5001    2202    PRI    PE1-PE6-         24320
Primary          UP    00:01:52  1 1 SE    -
10.12.183.6  13.1.1.1    5001    2202    DTR    PE1-PE6-         24321
Detour          UP    00:00:59  1 1 SE    -
Total 2 displayed, Up 2, Down 0.

```

```

PE1#sh rsvp session PE1-PE6-Detour
Ingress (Detour)
10.12.183.6
  From: 13.1.1.1, LSPstate: Up, LSPname: PE1-PE6-Detour
  Ingress FSM state: Operational
  Establishment Time: 0s 5ms
  SRLG-disjoint Configured: Preferred
  SRLG constraints are: 10
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled (Implicit), Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LDP Tunneling : Disabled
  LSP Protection: None
  Label in: -, Label out: 24321,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2202, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 13.1.1.2, xe8
  Path refresh: 30 seconds (RR enabled) (due in 29942 seconds)
  Resv lifetime: 157 seconds (due in 156 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Exclude path detail:
    Exclude Link: 12.1.1.2

```

```

Configured Path: none
Session Explicit Route Detail :
  13.1.1.2/32 strict
  25.1.1.1/32 strict
Record route:
-----
IP Address      Label
-----
<self>
13.1.1.2
25.1.1.1
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls

```

SRLG-disjoint as forced

```

PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint forced
PE1(config-router)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime     Rt   Style  Labelin  Labelout
10.12.183.6 10.12.183.1 5001     2202     PRI    PE1-PE6-         -
Primary      UP    00:00:03 1 1 SE    -      24320
10.12.183.6 13.1.1.1    5001     2202     DTR    PE1-PE6-         -
Detour      UP    00:00:03 1 1 SE    -      24321
Total 2 displayed, Up 2, Down 0.

PE1#sh rsvp session PE1-PE6-Detour
Ingress (Detour)
10.12.183.6
  From: 13.1.1.1, LSPstate: Up, LSPname: PE1-PE6-Detour
  Ingress FSM state: Operational
  Establishment Time: 0s 3ms
  SRLG-disjoint Configured: Forced
  SRLG constraints are: 10
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled (Implicit), Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LDP Tunneling : Disabled
  LSP Protection: None
  Label in: -, Label out: 24321,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2202, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 13.1.1.2, xe8
  Path refresh: 30 seconds (RR enabled) (due in 29988 seconds)
  Resv lifetime: 157 seconds (due in 138 seconds)
  Retry count: 0, intrvl: 30 seconds

```

```

RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Exclude path detail:
  Exclude Link: 12.1.1.2
Configured Path: none
Session Explicit Route Detail :
  13.1.1.2/32 strict
  25.1.1.1/32 strict
Record route:
-----
IP Address      Label
-----
<self>
13.1.1.2
25.1.1.1
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500

```

Secondary

```

PE1#sh running-config interface xe11
!
interface xe11
  load-interval 30
  ip address 12.1.1.1/24
  label-switching
  mpls traffic-eng srlg 10
  isis network point-to-point
  ip router isis ISIS-IGP-100
  enable-rsvp
!
PE1#sh running-config interface xe5
!
interface xe5
  load-interval 30
  ip address 40.1.1.1/24
  label-switching
  mpls traffic-eng srlg 10
  isis network point-to-point
  ip router isis ISIS-IGP-100
  enable-rsvp
!
rsvp-path PE1-PE6-via-PE2-PE4 mpls
  12.1.1.2 strict
  15.1.1.2 strict
  25.1.1.1 strict
!
!
rsvp-path PE1-PE6-Secondary mpls
  40.1.1.2 strict
  30.1.1.2 strict
!
!
rsvp-trunk PE1-PE6 ipv4
  reoptimize
  primary path PE1-PE6-via-PE2-PE4
  secondary path PE1-PE6-Secondary
  to 10.12.183.6
!

```

SRLG-disjoint as preferred

```
PE1#conf t
```

```

Enter configuration commands, one per line.  End with CNTL/Z.
PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint preferred
PE1(config-router)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime     Rt   Style  Labelin  Labelout
10.12.183.6 10.12.183.1 5001     2204     PRI    PE1-PE6-         24320
Primary                                UP      00:00:53 1 1 SE    -
10.12.183.6 10.12.183.1 5001     2205     SEC    PE1-PE6-         24320
Secondary                                UP      00:00:41 1 1 SE    -
Total 2 displayed, Up 2, Down 0.

PE1#
PE1#sh rsvp session PE1-PE6-Secondary
Ingress (Secondary)
10.12.183.6
  From: 10.12.183.1, LSPstate: Up, LSPname: PE1-PE6-Secondary
  Ingress FSM state: Operational
  Establishment Time: 0s 7ms
  SRLG-disjoint Configured: Preferred
  SRLG constraints are: 10
  SRLG configured in backup path are: 10
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LDP Tunneling : Disabled
  LSP Protection: None
  Label in: -, Label out: 24320,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2205, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 40.1.1.2, xe5
  Path refresh: 30 seconds (RR enabled) (due in 29919 seconds)
  Resv lifetime: 157 seconds (due in 156 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: secl (in use)
  Configured Explicit Route Detail :
    40.1.1.2/32 strict
    30.1.1.2/32 strict
  Session Explicit Route Detail :
    40.1.1.2/32 strict
    30.1.1.2/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  40.1.1.2
  30.1.1.2
  Style: Shared Explicit Filter

```

SRLG-disjoint as forced

```
PE1#conf t
```


Enter configuration commands, one per line. End with CNTL/Z.

```
PE1(config)#
PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint forced
PE1(config-router)#commit
```

PE1#sh rsvp session

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
 State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
 * indicates the session is active with local repair at one or more nodes
 (P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|---------------|----------|--------|------|----------|-------|
| Uptime Rt | Style Labelin | Labelout | | | | |
| 10.12.183.6 | 10.12.183.1 | 5001 | 2204 | PRI | PE1-PE6- | |
| Primary | UP | 00:04:05 | 1 1 SE | - | 24320 | |
| 10.12.183.6 | 10.12.183.1 | 5001 | 2205 | SEC | PE1-PE6- | |
| Secondary | DN | N/A | 0 0 SE | - | - | |

Total 2 displayed, Up 1, Down 1.

PE1#

-----srlg-disjoint preferred under router rsvp and srlg-disjoint preferred under rsvp-trunk-----

```
PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint preferred
PE1(config-router)#exit
PE1(config)#rsvp-trunk PE1-PE6 ipv4
PE1(config-trunk)#secondary srlg-disjoint preferred
PE1(config-trunk)#commit
```

PE1#sh rsvp session

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
 State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
 * indicates the session is active with local repair at one or more nodes
 (P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|---------------|----------|--------|------|----------|-------|
| Uptime Rt | Style Labelin | Labelout | | | | |
| 10.12.183.6 | 10.12.183.1 | 5001 | 2204 | PRI | PE1-PE6- | |
| Primary | UP | 00:05:39 | 1 1 SE | - | 24320 | |
| 10.12.183.6 | 10.12.183.1 | 5001 | 2205 | SEC | PE1-PE6- | |
| Secondary | UP | 00:00:04 | 1 1 SE | - | 24320 | |

Total 2 displayed, Up 2, Down 0.

-----srlg-disjoint preferred under router rsvp and srlg-disjoint forced under rsvp-trunk-----

```
PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint preferred
PE1(config-router)#exit
PE1(config)#rsvp-trunk PE1-PE6 ipv4
PE1(config-trunk)#secondary srlg-disjoint forced
PE1(config-trunk)#commit
```

PE1#sh rsvp session

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
 State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
 * indicates the session is active with local repair at one or more nodes
 (P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-----------|---------------|----------|--------|------|---------|-------|
| Uptime Rt | Style Labelin | Labelout | | | | |

```

10.12.183.6      10.12.183.1      5001      2201      SEC      PE1-PE6-
Secondary      DN      N/A      0 0 SE      -
10.12.183.6      10.12.183.1      5001      2204      PRI      PE1-PE6-
Primary      UP      00:00:03      1 1 SE      -      24320
Total 2 displayed, Up 1, Down 1.

```

PE1#

SRLG-disjoint forced under router rsvp and srlg-disjoint preferred under rsvp-trunk

```

PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint forced
PE1(config-router)#exit
PE1(config)#rsvp-trunk PE1-PE6 ipv4
PE1(config-trunk)#secondary srlg-disjoint preferred
PE1(config-trunk)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
10.12.183.6 10.12.183.1 5001      2204      PRI      PE1-PE6-
Primary      UP      00:05:39  1 1 SE      -      24320
10.12.183.6 10.12.183.1 5001      2205      SEC      PE1-PE6-
Secondary    UP      00:00:04  1 1 SE      -      24320
Total 2 displayed, Up 2, Down 0.

```

SRLG-disjoint forced under router rsvp and srlg-disjoint forced under rsvp-trunk

```

PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint forced
PE1(config-router)#exit
PE1(config)#rsvp-trunk PE1-PE6 ipv4
PE1(config-trunk)#secondary srlg-disjoint forced
PE1(config-trunk)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
10.12.183.6 10.12.183.1 5001      2201      SEC      PE1-PE6-
Secondary    DN      N/A      0 0 SE      -      -
10.12.183.6 10.12.183.1 5001      2204      PRI      PE1-PE6-
Primary      UP      00:00:03  1 1 SE      -      24320
Total 2 displayed, Up 1, Down 1.

PE1#

```

Secondary-priority

```

rsvp-path SP5 mpls
10.1.1.2 strict
20.1.1.1 strict
30.1.1.2 strict
!

```

```

rsvp-path SP4 mpls
  11.1.1.2 strict
  22.1.1.1 strict
!
rsvp-path SP3 mpls
  41.1.1.2 strict
  26.1.1.1 strict
!
rsvp-path SP2 mpls
  13.1.1.2 strict
  25.1.1.1 strict
!
rsvp-path SP1 mpls
  40.1.1.2 strict
  30.1.1.2 strict
!
rsvp-path PE1-PE6-via-PE2-PE4 mpls
  12.1.1.2 strict
  15.1.1.2 strict
  25.1.1.1 strict
!
!
!
rsvp-trunk PE1-PE6 ipv4
  reoptimize
  primary path PE1-PE6-via-PE2-PE4
  secondary-priority 1 path SP1
  secondary-priority 2 path SP2
  secondary-priority 3 path SP3
  secondary-priority 4 path SP4
  secondary-priority 5 path SP5
  to 10.12.183.6
!

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime    Rt   Style  Labelin Labelout
10.12.183.6  10.12.183.1  5001    2201     PRI    PE1-PE6-         24320
Primary    UP
10.12.183.6  10.12.183.1  5001    2207     SEC    PE1-PE6-Secondary-Priority-
1          UP    00:01:00  1 1 SE    -      24320
Total 2 displayed, Up 2, Down 0.

PE1#

PE1#sh running-config interface xe11
!
interface xe11
  load-interval 30
  ip address 12.1.1.1/24
  label-switching
  mpls traffic-eng srlg 10
  isis network point-to-point
  ip router isis ISIS-IGP-100
  enable-rsvp
!
!
PE1#sh running-config interface xe5
!
interface xe5
  load-interval 30

```

```

ip address 40.1.1.1/24
label-switching
mpls traffic-eng srlg 10
isis network point-to-point
ip router isis ISIS-IGP-100
enable-rsvp
!

PE1#sh rsvp session PE1-PE6-Secondary-Priority-1
Ingress (Secondary-Priority1)
10.12.183.6
  From: 10.12.183.1, LSPstate: Up, LSPname: PE1-PE6-Secondary-Priority-1
  Ingress FSM state: Operational
  Establishment Time: 0s 8ms
  SRLG configured in backup path are: 10
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LDP Tunneling : Disabled
  LSP Protection: None
  Label in: -, Label out: 24320,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2207, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 40.1.1.2, xe5
  Path refresh: 30 seconds (RR enabled) (due in 30007 seconds)
  Resv lifetime: 157 seconds (due in 152 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: SP1 (in use)
  Configured Explicit Route Detail :
    40.1.1.2/32 strict
    30.1.1.2/32 strict
  Session Explicit Route Detail :
    40.1.1.2/32 strict
    30.1.1.2/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  40.1.1.2
  30.1.1.2
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500

```

SRLG-disjoint as preferred

```

PE1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
PE1(config)#
PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint preferred
PE1(config-router)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

```

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-----------------|----------|--------|------|-----------------------------|-------|
| Uptime Rt | Style Labelin | Labelout | | | | |
| 10.12.183.6 | 10.12.183.1 | 5001 | 2201 | PRI | PE1-PE6- | |
| Primary | UP | 00:03:41 | 1 1 SE | - | 24320 | |
| 10.12.183.6 | 10.12.183.1 | 5001 | 2207 | SEC | PE1-PE6-Secondary-Priority- | |
| 1 UP | 00:00:02 1 1 SE | - | 24320 | | | |

Total 2 displayed, Up 2, Down 0.

PE1#

```

PE1#sh rsvp session PE1-PE6-Secondary-Priority-1
Ingress (Secondary-Priority1)
10.12.183.6
  From: 10.12.183.1, LSPstate: Up, LSPname: PE1-PE6-Secondary-Priority-1
  Ingress FSM state: Operational
  Establishment Time: 0s 9ms
  SRLG-disjoint Configured: Preferred
  SRLG constraints are: 10
  SRLG configured in backup path are: 10
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LDP Tunneling : Disabled
  LSP Protection: None
  Label in: -, Label out: 24320,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2207, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 40.1.1.2, xe5
  Path refresh: 30 seconds (RR enabled) (due in 29976 seconds)
  Resv lifetime: 157 seconds (due in 140 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: SP1 (in use)
  Configured Explicit Route Detail :
    40.1.1.2/32 strict
    30.1.1.2/32 strict
  Session Explicit Route Detail :
    40.1.1.2/32 strict
    30.1.1.2/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  40.1.1.2
  30.1.1.2
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500

```

SRLG-disjoint as forced

```

PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint forced
PE1(config-router)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

```
Ingress RSVP:
To      From      Tun-ID  LSP-ID  Type  LSPName      State
Uptime Rt  Style  Labelin Labelout
10.12.183.6  10.12.183.1  5001    2201    PRI    PE1-PE6-
Primary      UP    00:04:36  1 1 SE    -      24320
10.12.183.6  10.12.183.1  5001    2207    SEC    PE1-PE6-Secondary-Priority-
2    UP    00:00:10  1 1 SE    -      24321
Total 2 displayed, Up 2, Down 0.
```

```
PE1#sh rsvp session PE1-PE6-Secondary-Priority-2
```

```
Ingress (Secondary-Priority2)
```

```
10.12.183.6
```

```
From: 10.12.183.1, LSPstate: Up, LSPname: PE1-PE6-Secondary-Priority-2
Ingress FSM state: Operational
Establishment Time: 0s 7ms
SRLG-disjoint Configured: Forced
SRLG constraints are: 10
Setup priority: 7, Hold priority: 0, HOP limit: 255
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
IGP-Shortcut: Disabled, LSP metric: 20
LDP Tunneling : Disabled
LSP Protection: None
Label in: -, Label out: 24321,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5001, LSP Id: 2207, Ext-Tunnel Id: 10.12.183.1
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 13.1.1.2, xe8
Path refresh: 30 seconds (RR enabled) (due in 29991 seconds)
Resv lifetime: 157 seconds (due in 144 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: SP2 (in use)
Configured Explicit Route Detail :
 13.1.1.2/32 strict
 25.1.1.1/32 strict
Session Explicit Route Detail :
 13.1.1.2/32 strict
 25.1.1.1/32 strict
Record route:
-----
IP Address      Label
-----
<self>
13.1.1.2
25.1.1.1
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
```

SRLG-disjoint preferred under router rsvp and srlg-disjoint preferred under rsvp-trunk

```
PE1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint preferred
PE1(config-router)#exit
PE1(config)#rsvp-trunk PE1-PE6 ipv4
PE1(config-trunk)# secondary-priority 1 srlg-disjoint preferred
PE1(config-trunk)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
```

* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-------------|----------|---------|----------|-----------------------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 10.12.183.6 | 10.12.183.1 | 5001 | 2201 | PRI | PE1-PE6- | |
| Primary | UP | 00:13:02 | 1 1 SE | - | 24320 | |
| 10.12.183.6 | 10.12.183.1 | 5001 | 2207 | SEC | PE1-PE6-Secondary-Priority- | |
| 1 | UP | 00:00:04 | 1 1 SE | - | 24320 | |

Total 2 displayed, Up 2, Down 0.

PE1#

PE1#sh rsvp session PE1-PE6-Secondary-Priority-1

Ingress (Secondary-Priority1)

10.12.183.6

From: 10.12.183.1, LSPstate: Up, LSPname: PE1-PE6-Secondary-Priority-1

Ingress FSM state: Operational

Establishment Time: 0s 8ms

SRLG-disjoint Configured: Preferred

SRLG constraints are: 10

SRLG configured in backup path are: 10

Setup priority: 7, Hold priority: 0, HOP limit: 255

CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds

LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS

IGP-Shortcut: Disabled, LSP metric: 20

LDP Tunneling : Disabled

LSP Protection: None

Label in: -, Label out: 24320,

Tspec rate: 0, Fspec rate: 0

Policer: Not Configured

Tunnel Id: 5001, LSP Id: 2207, Ext-Tunnel Id: 10.12.183.1

Bind value: 0, Oper state: NA, Alloc mode: NA

Downstream: 40.1.1.2, xe5

Path refresh: 30 seconds (RR enabled) (due in 29991 seconds)

Resv lifetime: 157 seconds (due in 136 seconds)

Retry count: 0, intrvl: 30 seconds

RRO re-use as ERO: Disabled

Label Recording: Disabled

Admin Groups: none

Configured Path: SP1 (in use)

Configured Explicit Route Detail :

40.1.1.2/32 strict

30.1.1.2/32 strict

Session Explicit Route Detail :

40.1.1.2/32 strict

30.1.1.2/32 strict

Record route:

| IP Address | Label |
|------------|-------|
| <self> | |
| 40.1.1.2 | |
| 30.1.1.2 | |

Style: Shared Explicit Filter

Traffic type: controlled-load

Minimum Path MTU: 1500

Recorded Time : N/A

SRLG-disjoint preferred under router rsvp and srlg-disjoint forced under rsvp-trunk

PE1#conf t

Enter configuration commands, one per line. End with CNTL/Z.

PE1(config)#router rsvp

PE1(config-router)# srlg-disjoint preferred

PE1(config-router)#rsvp-trunk PE1-PE6 ipv4

```
% Warning: Executing the CLI from higher config mode level
PE1(config-trunk)#secondary-priority 1 srlg-disjoint forced
PE1(config-trunk)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To      From      Tun-ID  LSP-ID  Type  LSPName      State
Uptime  Rt  Style  Labelin Labelout
10.12.183.6  10.12.183.1  5001    2201    PRI    PE1-PE6-
Primary      UP    00:16:55  1 1 SE    -      24320
10.12.183.6  10.12.183.1  5001    2207    SEC    PE1-PE6-Secondary-Priority-
2    UP    00:00:23  1 1 SE    -      24321
Total 2 displayed, Up 2, Down 0.

PE1#

PE1#sh rsvp session PE1-PE6-Secondary-Priority-2
Ingress (Secondary-Priority2)
10.12.183.6
  From: 10.12.183.1, LSPstate: Up, LSPname: PE1-PE6-Secondary-Priority-2
  Ingress FSM state: Operational
  Establishment Time: 0s 7ms
  SRLG-disjoint Configured: Preferred
  SRLG constraints are: 10
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LDP Tunneling : Disabled
  LSP Protection: None
  Label in: -, Label out: 24321,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2207, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 13.1.1.2, xe8
  Path refresh: 30 seconds (RR enabled) (due in 29974 seconds)
  Resv lifetime: 157 seconds (due in 125 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: SP2 (in use)
  Configured Explicit Route Detail :
    13.1.1.2/32 strict
    25.1.1.1/32 strict
  Session Explicit Route Detail :
    13.1.1.2/32 strict
    25.1.1.1/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  13.1.1.2
  25.1.1.1
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
```


SRLG-disjoint forced under router rsvp and srlg-disjoint preferred under rsvp-trunk

```

PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint forced
PE1(config-router)#exit
PE1(config)#rsvp-trunk PE1-PE6 ipv4
PE1(config-trunk)#secondary-priority 1 srlg-disjoint preferred
PE1(config-trunk)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime    Rt    Style  Labelin  Labelout
10.12.183.6  10.12.183.1  5001      2201      PRI     PE1-PE6-         24320
Primary                                UP      00:00:03  1 1 SE  -
10.12.183.6  10.12.183.1  5001      2207      SEC     PE1-PE6-Secondary-Priority-
1    UP      00:00:03  1 1 SE  -      24320
Total 2 displayed, Up 2, Down 0.

PE1#

```

SRLG-disjoint forced under router rsvp and srlg-disjoint forced under rsvp-trunk

```

PE1(config)#router rsvp
PE1(config-router)#srlg-disjoint forced
PE1(config-router)#exit
PE1(config)#rsvp-trunk PE1-PE6 ipv4
PE1(config-trunk)#secondary-priority 1 srlg-disjoint forced
PE1(config-trunk)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime    Rt    Style  Labelin  Labelout
10.12.183.6  10.12.183.1  5001      2201      PRI     PE1-PE6-         24320
Primary                                UP      00:00:29  1 1 SE  -
10.12.183.6  10.12.183.1  5001      2207      SEC     PE1-PE6-Secondary-Priority-
2    UP      00:00:02  1 1 SE  -      24321
Total 2 displayed, Up 2, Down 0.

PE1#

```

Manual-Bypass**Configure the tunnel between PE1-PE6 on PE1 via PE1-PE2-PE4-PE6**

```

rsvp-path PE1-PE6-via-PE2-PE4 mpls
12.1.1.2 strict
15.1.1.2 strict
25.1.1.1 strict
!
!
!
rsvp-trunk PE1-PE6 ipv4

```

```

reoptimize
primary fast-reroute protection facility
primary fast-reroute node-protection
primary path PE1-PE6-via-PE2-PE4
to 10.12.183.6
!

```

Configure the rsvp-bypass sessions for link protection and node protection on PE2

```

rsvp-path link mpls
 16.1.1.2 strict
!
rsvp-path node mpls
 14.1.1.2 strict
 22.1.1.1 strict
!
!
!
!
rsvp-bypass link
 path link
 to 10.12.183.4
!
rsvp-bypass node
 path node
 to 10.12.183.6
!

PE2#sh rsvp session link-Bypass
Ingress (Bypass)
10.12.183.4
  From: 10.12.183.2, LSPstate: Up, LSPname: link-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 3ms
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 10
  LDP Tunneling : Disabled
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out: 3,
  Tspec rate: 0, Espec rate: 0
  Policer: Not Configured
  Tunnel Id: 5001, LSP Id: 2201, Ext-Tunnel Id: 10.12.183.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 16.1.1.2, cd42/3
  Path refresh: 30 seconds (RR enabled) (due in 29904 seconds)
  Resv lifetime: 157 seconds (due in 133 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: link (in use)
  Configured Explicit Route Detail :
    16.1.1.2/32 strict
  Session Explicit Route Detail :
    16.1.1.2/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  16.1.1.2
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
  Recorded Time : N/A

```

```

Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Total LSP protected : 0, Bandwidth in use : 0
PE2#sh rsvp session node-Bypass
Ingress (Bypass)
10.12.183.6
  From: 10.12.183.2, LSPstate: Up, LSPname: node-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 6ms
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Disabled, Re-Optimization Timer: NA, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LDP Tunneling : Disabled
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out: 24320,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5002, LSP Id: 2203, Ext-Tunnel Id: 10.12.183.2
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 14.1.1.2, xe46
  Path refresh: 30 seconds (RR enabled) (due in 29975 seconds)
  Resv lifetime: 157 seconds (due in 135 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: node (in use)
  Configured Explicit Route Detail :
    14.1.1.2/32 strict
    22.1.1.1/32 strict
  Session Explicit Route Detail :
    14.1.1.2/32 strict
    22.1.1.1/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  14.1.1.2
  22.1.1.1
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Total LSP protected : 1, Bandwidth in use : 0
PE2#

```

Verify whether the bypass link is protecting the Tunnel

```

Verify whether the bypass link is protecting the Tunnel
-----

PE2#sh rsvp bypass protected-lsp-list

```

Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node protection miss,
 3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: link
 Bypass trunk bandwidth type: best-effort
 Total LSP protected : 0
 Total LSPs using protection path : 0
 Bandwidth in use : 0

Bypass trunk: node
 Bypass trunk bandwidth type: best-effort
 List of LSP's Protected:

| Tunnel-id | Lsp-Id | Lsp-Name | Role | Ext_tnl_ |
|-----------|---------|-----------------|------------|-------------|
| id | Ingress | Egress | Match-Code | Mapped-BW |
| 5001 | 2201 | PE1-PE6-Primary | Transit | 10.12.183.1 |
| | | | | 10.12.183.1 |
| | | | | 10.12.183.6 |
| | | | | 0 |
| | | | | NA |

Total LSP protected : 1
 Total LSPs using protection path : 0
 Bandwidth in use : 0

PE2#

Configure the srlg values on primary and node bypass secondary path and configure the srlg-disjoint preferred under router rsvp

```
PE2#sh running-config interface cd42/1
!
interface cd42/1
 load-interval 30
 ip address 15.1.1.1/24
 label-switching
 isis network point-to-point
 ip router isis ISIS-IGP-100
 enable-rsvp
!
!
PE2#sh running-config interface cd42/3
!
interface cd42/3
 load-interval 30
 ip address 16.1.1.1/24
 label-switching
 isis network point-to-point
 ip router isis ISIS-IGP-100
 enable-rsvp
!
!
PE2#sh running-config interface xe46
!
interface xe46
 load-interval 30
 ip address 14.1.1.1/24
 label-switching
 isis network point-to-point
 ip router isis ISIS-IGP-100
 enable-rsvp
!
!
PE2#
```

```
PE2(config)#int cd42/1
PE2(config-if)#mpls traffic-eng srlg 5
PE2(config-if)#interface xe46
PE2(config-if)#mpls traffic-eng srlg 5
PE2(config-if)#commit
```

Configure the srlg-disjoint preferred under router rsvp

```

PE2(config-if)#router rsvp
PE2(config-router)#srlg-disjoint preferred
PE2(config-router)#commit

PE2#
PE2#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime    Rt  Style  Labelin Labelout
10.12.183.4  10.12.183.2  5001    2201    BPS    link-            3
Bypass                    UP      00:03:53  1 1 SE    -                24320
10.12.183.6  10.12.183.2  5002    2203    BPS    node-            24320
Bypass                    UP      00:02:48  1 1 SE    -                24320
Total 2 displayed, Up 2, Down 0.

Transit RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime    Rt  Style  Labelin Labelout
10.12.183.6  10.12.183.1  5001    2201    PRI    PE1-PE6-        24320
Primary                    UP      00:09:09  1 1 SE    24320
Total 1 displayed, Up 1, Down 0.

PE2#
PE2#sh rsvp bypass protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
              3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: link
Bypass trunk bandwidth type: best-effort
Total LSP protected : 0
Total LSPs using protection path : 0
Bandwidth in use : 0

Bypass trunk: node
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id  Lsp-Id   Lsp-Name          Role          Ext_tnl_
id         Ingress  Egress            Match-Code    Mapped-BW
5001      2201     PE1-PE6-          10.12.183.1   10.12.183.1   10.12.183.6   0           NA
Primary    Transit
Total LSP protected : 1
Total LSPs using protection path : 0
Bandwidth in use : 0

PE2#

```

Configure the srlg-disjoint forced under router rsvp

```

PE2(config)#router rsvp
PE2(config-router)#srlg-disjoint forced
PE2(config-router)#commit

PE2#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes

```

(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-------------|--------|----------|----------|---------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 10.12.183.4 | 10.12.183.2 | 5001 | 2201 | BPS | link- | |
| Bypass | | UP | 00:03:20 | 1 1 SE | - | 3 |
| 10.12.183.6 | 10.12.183.2 | 5002 | 2203 | BPS | node- | |
| Bypass | | UP | 00:03:20 | 1 1 SE | - | 24320 |

Total 2 displayed, Up 2, Down 0.

Transit RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-------------|--------|----------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 10.12.183.6 | 10.12.183.1 | 5001 | 2201 | PRI | PE1-PE6- | |
| Primary | | UP | 00:03:16 | 1 1 SE | 24320 | 24320 |

Total 1 displayed, Up 1, Down 0.

PE2#sh rsvp bypass protected-lsp-list

Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node protection miss,
3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: link

Bypass trunk bandwidth type: best-effort

List of LSP's Protected:

| Tunnel-id | Lsp-Id | Lsp-Name | Role | Ext_tnl_ |
|-----------|---------|----------|-------------|-------------|
| id | Ingress | Egress | Match-Code | Mapped-BW |
| 5001 | 2201 | PE1-PE6- | | |
| Primary | | Transit | 10.12.183.1 | 10.12.183.1 |
| | | | | 10.12.183.6 |
| | | | 2 | NA |

Total LSP protected : 1

Total LSPs using protection path : 0

Bandwidth in use : 0

Bypass trunk: node

Bypass trunk bandwidth type: best-effort

Total LSP protected : 0

Total LSPs using protection path : 0

Bandwidth in use : 0

PE2#

RSVP Auto-bypass

```
router rsvp
  auto-bypass
  attributes best-effort
  reoptimize
  exit
  inactivity-timer 1
  enable
  exit

rsvp-path PE1-PE6-via-PE2-PE4 mpls
  12.1.1.2 strict
  15.1.1.2 strict
  25.1.1.1 strict
!
!
!
rsvp-trunk PE1-PE6 ipv4
  reoptimize
  primary fast-reroute protection facility
  primary fast-reroute node-protection
  primary path PE1-PE6-via-PE2-PE4
  to 10.12.183.6
!
```



```

25.1.1.1/32 strict
Record route:
LP = 1 -> PLR's Downstream link is protected      PU = 1 -> Protection is in use on PLR
NP = 1 -> PLR's Downstream neighbor is protected  BP = 1 -> BW protection available at PLR
-----
IP Address      Label      (LP, PU, NP, BP)
-----
<self>
12.1.1.2        24320      ( 1,  0,  1,  0)
15.1.1.2        24320      ( 0,  0,  0,  0)
25.1.1.1        24320      ( 0,  0,  0,  0)
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
PE1#sh rsvp session BN-10.12.183.2-10.12.183.4-101-Bypass
Ingress (Bypass)
10.12.183.4
  From: 10.12.183.1, LSPstate: Up, LSPname: BN-10.12.183.2-10.12.183.4-101-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 6ms
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  Auto-bypass Info: Group-ID 101, Exclude-Node 10.12.183.2, Egress 10.12.183.4, Cspf ISIS
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 10
  LDP Tunneling : Disabled
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out:      3,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5003, LSP Id: 2205, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 41.1.1.2, xe9
  Path refresh: 30 seconds (RR enabled) (due in 29890 seconds)
  Resv lifetime: 157 seconds (due in 150 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: none
  Exclude path detail:
    Exclude Node: 10.12.183.2
  Session Explicit Route Detail :
    41.1.1.2/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  41.1.1.2
  Style: Shared Explicit Filter
  Traffic type: controlled-load
  Minimum Path MTU: 1500
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None

```



```

    Originated Node : None, Recorded Time : N/A
    Trunk Type: mpls
    Total LSP protected : 1, Bandwidth in use : 0
PE1#

PE1#sh rsvp bypass protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
           3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: BL-12.1.1.2-10.12.183.2-101
Bypass trunk bandwidth type: best-effort
Total LSP protected : 0 (inactivity timeout due in 47 seconds)
Total LSPs using protection path : 0
    Bandwidth in use : 0

Bypass trunk: BN-10.12.183.2-10.12.183.4-101
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id  Lsp-Id    Lsp-Name          Role          Ext_tnl_
id         Ingress    Egress            Match-Code    Mapped-BW
5001      2201      PE1-PE6-          10.12.183.1   10.12.183.1   10.12.183.6     0      NA
Primary    Ingress
Total LSP protected : 1
Total LSPs using protection path : 0
    Bandwidth in use : 0

PE1#sh rsvp auto-bypass-protected-lsp-list
Tunnel-id  Lsp-Id    Lsp-Name          Role          Ext_tnl_
id         Ingress    Egress            Bypass-Name
5001      2201      PE1-PE6-          10.12.183.1   10.12.183.1   10.12.183.6     BN-
Primary    Ingress    10.12.183.2-10.12.183.4-101
PE1#

```

Configure the srlg values on both primary and auto-bypass path

```

PE1#sh running-config interface xe11
!
interface xe11
 load-interval 30
 ip address 12.1.1.1/24
 label-switching
 isis network point-to-point
 ip router isis ISIS-IGP-100
 enable-rsvp
!

PE1#sh running-config interface xe9
!
interface xe9
 load-interval 30
 ip address 41.1.1.1/24
 label-switching
 isis network point-to-point
 ip router isis ISIS-IGP-100
 enable-rsvp
!

PE1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
PE1(config)#int xe11
PE1(config-if)# mpls traffic-eng srlg 10
PE1(config-if)#int xe9
% Warning: Executing the CLI from higher config mode level
PE1(config-if)# mpls traffic-eng srlg 10
PE1(config-if)#commit

```

Configure srlg-disjoint as preferred under router rsvp

```

PE1(config-if)#router rsvp
% Warning: Executing the CLI from higher config mode level
PE1(config-router)#srlg-disjoint preferred
PE1(config-router)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID   LSP-ID   Type   LSPName          State
Uptime    Rt   Style  Labelin  Labelout
10.12.183.6  10.12.183.1  5001    2201    PRI    PE1-PE6-         UP
Primary                                24320
10.12.183.4  10.12.183.1  5002    2204    BPS    BN-10.12.183.2-10.12.183.4-100-Bypass
UP 00:00:
02 1 1 SE    -          3
10.12.183.2  10.12.183.1  5003    2205    BPS    BL-12.1.1.2-10.12.183.2-100-Bypass
UP 00:00:
02 1 1 SE    -          24321
Total 3 displayed, Up 3, Down 0.

PE1#

PE1#sh rsvp session BN-10.12.183.2-10.12.183.4-100-Bypass
Ingress (Bypass)
10.12.183.4
  From: 10.12.183.1, LSPstate: Up, LSPname: BN-10.12.183.2-10.12.183.4-100-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 4ms
  SRLG-disjoint Configured: Preferred
  SRLG constraints are: 10
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  Auto-bypass Info: Group-ID 100, Exclude-Node 10.12.183.2, Egress 10.12.183.4, Cspf ISIS
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 10
  LDP Tunneling : Disabled
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out: 3,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5002, LSP Id: 2204, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 13.1.1.2, xe8
  Path refresh: 30 seconds (RR enabled) (due in 29974 seconds)
  Resv lifetime: 157 seconds (due in 140 seconds)
  Retry count: 0, intrvl: 30 seconds
  RRO re-use as ERO: Disabled
  Label Recording: Disabled
  Admin Groups: none
  Configured Path: none
  Exclude path detail:
    Exclude Node: 10.12.183.2
  Session Explicit Route Detail :
    13.1.1.2/32 strict
  Record route:
  -----
  IP Address      Label
  -----
  <self>
  13.1.1.2
  Style: Shared Explicit Filter

```

```
Traffic type: controlled-load
Minimum Path MTU: 1500

PE1#sh rsvp bypass protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
          3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

Bypass trunk: BN-10.12.183.2-10.12.183.4-100
Bypass trunk bandwidth type: best-effort
List of LSP's Protected:
Tunnel-id Lsp-Id Lsp-Name Role Ext_tnl_
id Ingress Egress Match-Code Mapped-BW
5001 2201 PE1-PE6- Primary Ingress 10.12.183.1 10.12.183.1 10.12.183.6 0 NA
Total LSP protected : 1
Total LSPs using protection path : 0
Bandwidth in use : 0

Bypass trunk: BL-12.1.1.2-10.12.183.2-100
Bypass trunk bandwidth type: best-effort
Total LSP protected : 0 (inactivity timeout due in 10 seconds)
Total LSPs using protection path : 0
Bandwidth in use : 0

PE1#sh rsvp auto-bypass-protected-lsp-list
Tunnel-id Lsp-Id Lsp-Name Role Ext_tnl_
id Ingress Egress Bypass-Name
5001 2201 PE1-PE6- Primary Ingress 10.12.183.1 10.12.183.1 10.12.183.6 BN-
10.12.183.2-10.12.183.4-100
PE1#
```

Configure srlg-disjoint as forced under router rsvp

```
PE1(config-if)#router rsvp
% Warning: Executing the CLI from higher config mode level
PE1(config-router)#srlg-disjoint forced
PE1(config-router)#commit

PE1#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To From Tun-ID LSP-ID Type LSPName State
Uptime Rt Style Labelin Labelout
10.12.183.6 10.12.183.1 5001 2201 PRI PE1-PE6-
Primary UP 00:12:20 1 1 SE - 24320
10.12.183.4 10.12.183.1 5002 2204 BPS BN-10.12.183.2-10.12.183.4-100-Bypass
UP 00:00:
04 1 1 SE - 3
10.12.183.2 10.12.183.1 5003 2205 BPS BL-12.1.1.2-10.12.183.2-100-Bypass
UP 00:00:
04 1 1 SE - 24321
Total 3 displayed, Up 3, Down 0.

PE1#sh rsvp session BN-10.12.183.2-10.12.183.4-100-Bypass
Ingress (Bypass)
10.12.183.4
From: 10.12.183.1, LSPstate: Up, LSPname: BN-10.12.183.2-10.12.183.4-100-Bypass
Ingress FSM state: Operational
Establishment Time: 0s 2ms
SRLG-disjoint Configured: Forced
```

```

SRLG constraints are: 10
Setup priority: 7, Hold priority: 0, HOP limit: 255
Auto-bypass Info: Group-ID 100, Exclude-Node 10.12.183.2, Egress 10.12.183.4, Cspf ISIS
CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
IGP-Shortcut: Disabled, LSP metric: 10
LDP Tunneling : Disabled
LSP Protection: None
Bypass trunk bandwidth type: Best-effort
  Label in: -, Label out:      3,
Tspec rate: 0, Fspec rate: 0
Policer: Not Configured
Tunnel Id: 5002, LSP Id: 2204, Ext-Tunnel Id: 10.12.183.1
Bind value: 0, Oper state: NA, Alloc mode: NA
Downstream: 13.1.1.2, xe8
Path refresh: 30 seconds (RR enabled) (due in 29999 seconds)
Resv lifetime: 157 seconds (due in 147 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: none
Exclude path detail:
  Exclude Node: 10.12.183.2
Session Explicit Route Detail :
  13.1.1.2/32 strict
Record route:
-----
IP Address      Label
-----
<self>
13.1.1.2
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500
Recorded Time : N/A
Current Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Last Signaled Error:
  Code : None, Value : None
  Originated Node : None, Recorded Time : N/A
Trunk Type: mpls
Total LSP protected : 1, Bandwidth in use : 0
PE1#

PE1#sh rsvp session BL-12.1.1.2-10.12.183.2-100-Bypass
Ingress (Bypass)
10.12.183.2
  From: 10.12.183.1, LSPstate: Up, LSPname: BL-12.1.1.2-10.12.183.2-100-Bypass
  Ingress FSM state: Operational
  Establishment Time: 0s 6ms
  SRLG-disjoint Configured: Forced
  SRLG constraints are: 10
  Setup priority: 7, Hold priority: 0, HOP limit: 255
  Auto-bypass Info: Group-ID 100, Exclude-Link 12.1.1.2, Egress 10.12.183.2, Cspf ISIS
  CSPF usage: Enabled, CSPF Retry Count: 0, CSPF Retry Interval: 30 seconds
  LSP Re-Optimization: Enabled, Re-Optimization Timer: 5 minutes, Cspf Client: ISIS
  IGP-Shortcut: Disabled, LSP metric: 20
  LDP Tunneling : Disabled
  LSP Protection: None
  Bypass trunk bandwidth type: Best-effort
    Label in: -, Label out:    24321,
  Tspec rate: 0, Fspec rate: 0
  Policer: Not Configured
  Tunnel Id: 5003, LSP Id: 2205, Ext-Tunnel Id: 10.12.183.1
  Bind value: 0, Oper state: NA, Alloc mode: NA
  Downstream: 13.1.1.2, xe8

```

```

Path refresh: 30 seconds (RR enabled) (due in 29959 seconds)
Resv lifetime: 157 seconds (due in 128 seconds)
Retry count: 0, intrvl: 30 seconds
RRO re-use as ERO: Disabled
Label Recording: Disabled
Admin Groups: none
Configured Path: none
Exclude path detail:
  Exclude Link: 12.1.1.2
Session Explicit Route Detail :
  13.1.1.2/32 strict
  15.1.1.1/32 strict
Record route:
-----
IP Address      Label
-----
<self>
13.1.1.2
15.1.1.1
Style: Shared Explicit Filter
Traffic type: controlled-load
Minimum Path MTU: 1500

```

```

PE1#sh rsvp bypass protected-lsp-list
Match Code: 0 - Perfect match (all criteria matching), 1 - Bandwidth protection miss, 2 - Node
protection miss,
           3 - SRLG protection miss, 4 - Merge point not ideal, 255 - Invalid

```

```
Bypass trunk: BN-10.12.183.2-10.12.183.4-100
```

```
Bypass trunk bandwidth type: best-effort
```

```
List of LSP's Protected:
```

| Tunnel-id | Lsp-Id | Lsp-Name | Role | Ext_tnl_ |
|--------------------------------------|---------|-------------|----------------------|------------------|
| id | Ingress | Egress | Match-Code Mapped-BW | |
| 5001 | 2201 | PE1-PE6- | | |
| Primary | Ingress | 10.12.183.1 | 10.12.183.1 | 10.12.183.6 0 NA |
| Total LSP protected : 1 | | | | |
| Total LSPs using protection path : 0 | | | | |
| Bandwidth in use : 0 | | | | |

```
Bypass trunk: BL-12.1.1.2-10.12.183.2-100
```

```
Bypass trunk bandwidth type: best-effort
```

```
Total LSP protected : 0 (inactivity timeout due in 15 seconds)
```

```
Total LSPs using protection path : 0
```

```
Bandwidth in use : 0
```

```
PE1#sh rsvp auto-bypass-protected-lsp-list
```

| Tunnel-id | Lsp-Id | Lsp-Name | Role | Ext_tnl_ |
|-----------------------------|---------|-------------|-------------|-----------------|
| id | Ingress | Egress | Bypass-Name | |
| 5001 | 2201 | PE1-PE6- | | |
| Primary | Ingress | 10.12.183.1 | 10.12.183.1 | 10.12.183.6 BN- |
| 10.12.183.2-10.12.183.4-100 | | | | |

```
PE1#
```

SRLG with ISIS IGP Commands

The following commands are applicable for SRLG with ISIS IGP commands:

- [mpls traffic-eng srlg \(page 1425\)](#)
- [secondary-priority srlg-disjoint \(page 1434\)](#)
- [secondary srlg-disjoint \(page 1435\)](#)
- [srlg-disjoint \(page 1483\)](#)

Implementation Examples

SRLG with One-to-One FRR Protection

Scenario 1: When SRLG is configured only on Primary and Detour LSP's

Usecase 1: When we are configuring the SRLG value on Primary LSP and the Detour LSP and configured srlg-disjoint as preferred/forced under the router RSVP, then the detour LSP with same SRLG value will go down and new detour LSP will come UP where SRLG is not matching with the primary LSP SRLG value.

Scenario 2: When SRLG is configured on Primary and also on all available paths.

Usecase 2: When we are configuring the SRLG value on Primary LSP and also on all available paths and configured srlg-disjoint preferred under the router RSVP, then the detour session will be UP with the best path. The srlg-disjoint forced under the router RSVP, then the detour session will be DOWN as there will be not available path which is not matching the primary LSP SRLG value.

SRLG with Secondary Protection

Scenario 1: When SRLG is configured only on Primary and dynamic Secondary LSP's and configuring srlg-disjoint preferred under router RSVP.

Usecase 1: When we are configuring the SRLG value on Primary LSP and the secondary LSP and configured srlg-disjoint as preferred/forced under the router RSVP, then the secondary LSP with same SRLG value will go down and new secondary LSP will come UP where SRLG is not matching with the primary LSP SRLG value.

Scenario 2: When SRLG is configured on Primary and also on all available paths and configuring srlg-disjoint preferred under router RSVP.

Usecase 2: When we are configuring the SRLG value on Primary LSP and also on all available paths and configured srlg-disjoint preferred under the router RSVP, then the secondary session will be UP with the best path. The srlg-disjoint forced under the router RSVP, then the secondary session will be DOWN as there will be not available path which is not matching the primary LSP SRLG value.

Scenario 3: When SRLG is configured on Primary and defined Secondary LSP's and configuring srlg-disjoint preferred under router RSVP.

Usecase 3: When we are configuring the SRLG value on Primary LSP and also on the defined secondary LSP and configured srlg-disjoint preferred under the router RSVP, then the secondary session will be UP with the defined path. srlg-disjoint forced under the router RSVP, then the secondary session will be DOWN.

Scenario 4: When SRLG is configured only on Primary and dynamic Secondary LSP's and configuring srlg-disjoint preferred under router rsvp and srlg-disjoint preferred/forced under rsvp-trunk.

Usecase 4: When we are configuring the SRLG value on Primary LSP and the secondary LSP and configured srlg-disjoint as preferred/forced under the router RSVP and srlg-disjoint preferred/forced under rsvp-trunk, then the secondary LSP with same SRLG value will go down and new secondary LSP will come UP where SRLG is not matching with the primary LSP SRLG value.

Scenario 5: When SRLG is configured on Primary and also on all available paths and configuring srlg-disjoint preferred under router RSVP and srlg-disjoint preferred/forced under rsvp-trunk..

Usecase 5: When we are configuring the SRLG value on Primary LSP and also on all available paths and configured srlg-disjoint preferred under the router RSVP and srlg-disjoint preferred under RSVP trunk, then the secondary session will be UP with the best path. srlg-disjoint preferred under the router RSVP and srlg-disjoint forced under RSVP trunk, then the secondary session will be DOWN as there will be not available path which is not matching the primary LSP SRLG value.

Scenario 6: When SRLG is configured on Primary and defined Secondary LSP's and configuring srlg-disjoint forced under router RSVP and srlg-disjoint preferred/forced under rsvp-trunk.

Usecase 6: When we are configuring the SRLG value on Primary LSP and also on the defined secondary LSP and configured srlg-disjoint forced under the router RSVP and srlg-disjoint preferred under RSVPtrunk, then the secondary session will be UP with the defined path. The srlg-disjoint forced under the router RSVP and srlg-disjoint forced under RSVPtrunk, then the secondary session will be DOWN.

SRLG with Secondary-Priority Protection

Scenario 1: When SRLG is configured on Primary and on Secondary-priority path 1 and configuring srlg-disjoint forced under router rsvp.

Usecase 1: Configure five secondary-priority paths, and then attach all SP's under the configured rsvp trunk. When we are configuring the SRLG value on Primary LSP and on the secondary-priority 1(SP1) and configured srlg-disjoint preferred under the router rsvp, then the SP1 will be UP. The srlg-disjoint forced under the router rsvp, then the SP1 will go down and the SP2 comes UP.

Scenario 2: When SRLG is configured on Primary and on Secondary-priority path 1 and configuring srlg-disjoint preferred under router RSVP and srlg-disjoint preferred/forced under rsvp-trunk.

Usecase 2: Configure five secondary-priority paths, and then attach all SP's under the configured RSVP trunk. When we are configuring the SRLG value on Primary LSP and on the secondary-priority 1 (SP1) and configured srlg-disjoint preferred under the router RSVP and srlg-disjoint preferred under RSVP trunk, then the SP1 will be UP. The srlg-disjoint preferred under the router RSVP and srlg-disjoint forced under RSVP trunk, then the SP2 will be UP.

Scenario-3: When SRLG is configured on Primary and on Secondary-priority path 1 and configuring srlg-disjoint forced under router RSVP and srlg-disjoint preferred/forced under rsvp-trunk.

Usecase 3: Configure five secondary-priority paths, and then attach all SP's under the configured RSVP trunk. When we are configuring the SRLG value on Primary LSP and on the secondary-priority 1(SP1) and configured srlg-disjoint forced under the router RSVP and srlg-disjoint preferred under RSVP trunk, then the SP1 will be UP. The srlg-disjoint forced under the router RSVP and srlg-disjoint forced under RSVP trunk, then the SP2 will be UP.



Note: The SRLG option configured on the tunnel will take higher preference to the one configured on router-rsvp instance.

SRLG with Manual Bypass Protection

Scenario 1: When SRLG is configured on Primary and defined Manual bypass path and configuring srlg-disjoint preferred/forced under router RSVP.

Usecase 1: When we are configuring the SRLG value on Primary LSP and also on the defined manual bypass LSP and configured srlg-disjoint preferred under the router RSVP, then the bypass session will be UP and protecting the primary path. The srlg-disjoint forced under the router RSVP, then the bypass session will be UP but will not be protecting the primary path.

Scenario 2: When SRLG is configured on Primary and dynamic Manual bypass path and configuring srlg-disjoint preferred/forced under router RSVP.

Usecase 2: When we are configuring the SRLG value on Primary LSP and also on the dynamic manual bypass LSP and configured srlg-disjoint preferred under the router RSVP, then the bypass session will be UP and protecting the primary path. The srlg-disjoint forced under the router RSVP, then the bypass session will be UP but will not be protecting the primary path.

SRLG with Auto Bypass Protection

Scenario 1: When SRLG is configured on Primary and on Auto bypass path and configuring srlg-disjoint preferred/forced under router RSVP.

Usecase 1: When we are configuring the SRLG value on Primary LSP and also on the auto bypass LSP and configured srlg-disjoint preferred under the router RSVP, then the auto bypass session formed will change the path with the other available paths. The srlg-disjoint forced under the router RSVP, then the auto bypass session formed will change the path with the other available paths.

Scenario 2: When SRLG is configured on Primary and also on all available paths and configuring srlg-disjoint preferred/forced under router RSVP.

Usecase 2: When we are configuring the SRLG value on Primary LSP and also on all other LSP and configured srlg-disjoint preferred under the router RSVP, then the auto bypass session will be formed with the best path. The srlg-disjoint forced under the router RSVP, then the auto bypass session will go down since there will be no paths which is not matching with the primary LSP SRLG values.

Troubleshooting

1. When SRLG value is configured on both primary and on detour/secondary/bypass path and observing that the detour/secondary/bypass session is still considering the same SRLG value configured link, then check for the srlg-disjoint config is configured or not under router RSVP or rsvp-trunk.
2. In case of secondary/secondary-priority case, when you configured srlg-disjoint as preferred under router RSVP but you observe the srlg-disjoint forced behaviour then we need to check the configuration under the rsvp-trunk whether srlg-disjoint forced is configured.
3. When you configure the SRLG value on the interfaces apart from primary and you still see the detour/secondary/bypass path, then once cross verify the SRLG value is configured under the downstream interface or upstream interface.
4. When the RSVP session are UP with the SRLG constraints and you change the SRLG value on primary/backup path and you don't see the SRLG values are not considered immediately. Then you need to check for the reoptimization timer, because the modification of the SRLG values will take effect only after the expiry of the reoptimization timer.
5. When you modify the SRLG values and you observe that the SRLG values are not considered even after the reoptimization timer, then we need to check whether the primary session is UP or DOWN, because whatever the SRLG value modifications done when the primary is down/not in use, then the modified srlg values will be considered only when the primary session comes UP and the reoptimization timer expires.
6. In case of manual bypass, when manual bypass is configured and session is UP and protecting the primary path, then you modify the SRLG value on the bypass path, but we find even after the reoptimization timer the session is still up, the behaviour is expected and we need to check that the bypass tunnel will be UP and will not be protecting the primary trunk.
7. When we modify the SRLG values on primary and non primary links and we find the SRLG values are not considered and changes the paths, then we need to check whether the reoptimization is configured or not. The modifications of the SRLG values works only when the reoptimization is configured.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms/Acronym | Description |
|-------------------|---|
| RSVP-TE | Resource Reservation Protocol for Traffic Engineering; sets up TE LSPs. |
| MPLS | Multi-Protocol Label Switching; forwards packets using labels instead of IP lookup. |
| IGP | Interior Gateway Protocol (ISIS used here). |
| ISIS L1 | Level-1 ISIS routing within a single area. |
| CSPF | Constraint-Based Shortest Path First; used for TE LSP calculation. |
| FRR | Fast Reroute; fast protection switching for LSPs. |
| SRLG | Shared Risk Link Group; identifies links that may fail together. |
| LSDB | Link State Database; used by ISIS/OSPF to compute routes. |
| TE Link | Link with MPLS TE and RSVP enabled. |
| LSP | Label-Switched Path, established via RSVP-TE. |
| BFD | Bidirectional Forwarding Detection; rapid link failure detection. |

MPLS DiffServ Configuration

This chapter contains an overview of MPLS DiffServ functionality and terminology, MPLS DiffServ configuration example for a relevant scenario, configuration guidelines, and sample procedures for enabling and configuring MPLS DiffServ.

MPLS Diff-Serv Overview

The initial efforts to provide quality of service (QoS) in IP networks were based on a per application-Flow model (IntServ), in which individual applications requested QoS. With large number of flows traversing IP networks, this approach proved to be un-scalable and overly complex, and a more “coarse-grained” model was developed in the form of DiffServ. DiffServ approaches the problem of QoS by dividing traffic into a small number of classes and allocating network resources on a per-class basis. DiffServ provides differential forwarding treatment to traffic, thus enforcing QoS for different traffic flows. It is a scalable solution that does not require per flow signalling and state maintenance in the core. However, it cannot guarantee QoS if the path followed by the traffic does not have adequate resources to meet the QoS requirements.

DiffServ Tunnelling modes:

RFC 3270 has recommended three QoS models for DiffServ tunnelled traffic in MPLS networks:

OcNOS supports two models:

- Pipe model (default mode): With the Pipe Model, MPLS tunnels (aka LSPs) are used to hide the intermediate MPLS nodes between LSP Ingress and Egress from the Diff-Serv perspective. In this model, tunneled packets must convey two meaningful pieces of Diff-Serv information:
 - The Diff-Serv information which is meaningful to intermediate nodes along the LSP span including the LSP Egress (which we refer to as the “LSP Diff-Serv Information”). This LSP Diff-Serv Information is not meaningful beyond the LSP Egress: Whether Traffic Conditioning at intermediate nodes on the LSP span affects the LSP Diff-Serv information or not, this updated Diff-Serv information is not considered meaningful beyond the LSP Egress and is ignored.
 - The Diff-Serv information which is meaningful beyond the LSP Egress (which we refer to as the “Tunneled Diff-Serv Information”). This information is to be conveyed by the LSP Ingress to the LSP Egress. This Diff-Serv information is not meaningful to the intermediate nodes on the LSP span.
- Uniform model: With the Uniform Model, MPLS tunnels (aka LSPs) are viewed as artifacts of the end-to-end path from the Diff-Serv standpoint. MPLS Tunnels may be used for forwarding purposes but have no significant impact on Diff-Serv. In this model, any packet contains exactly one piece of Diff-Serv information which is meaningful and is always encoded in the outer most label entry (or in the IP DSCP where the IP packet is transmitted unlabeled for instance at the egress of the LSP). Any Diff-Serv information encoded somewhere else (e.g., in deeper label entries) is of no significance to intermediate nodes or to the tunnel egress and is ignored. If Traffic Conditioning at intermediate nodes on the LSP span affects the “outer” Diff-Serv information, the updated Diff-Serv information is the one considered meaningful at the egress of the LSP.
 - The Uniform Model for Diff-Serv over MPLS is such that, from the Diff-Serv perspective, operations are exactly identical to the operations if MPLS was not used. In other words, MPLS is entirely transparent to the Diff-Serv operations.
 - Use of the Uniform Model allows LSPs to span Diff-Serv domain boundaries without any other measure in place than an inter-domain Traffic Conditioning Agreement at the physical boundary

between the Diff-Serv domains and operating exclusively on the “outer” header, since the meaningful Diff-Serv information is always visible and modifiable in the outmost label entry.

Terminology

Following is a brief description of terms and concepts used to describe MPLS Diffserv.

EXP Value

The MPLS experimental bits (EXP) field is a 3-bit field in the MPLS header that you can use to define the QoS treatment (per-hop behavior) that a node should give to a packet. In an IP network, the DiffServ Code Point (DSCP) (a 6-bit field) defines a class and drop precedence. The EXP bits can be used to carry some of the information encoded in the IP DSCP and can also be used to encode the dropping precedence.

By default, OcNOS copies the three most significant bits of the DSCP or the IP precedence of the IP packet to the EXP field in the MPLS header. This action happens when the MPLS header is initially imposed on the IP packet. However, you can also set the EXP field by defining a mapping between the DSCP or IP precedence and the EXP bits. This mapping is configured using the `set mpls class` command in `pmap-class` mode or `qos map class exp` in global mode. For more information, see the “Remarking” section.

DSCP Value

Differentiated Services Code Point (DSCP) is a 6-bit value used to classify the priority of Layer-3 packets upon entry into a network. DSCP values range from 0 to 63, 63 being the highest priority, 0 being best-effort traffic.

Classification

Traffic classification allows the network to recognize traffic as it falls into classes that you have configured. Network traffic must be classified to apply specific QoS to it. Classification can be inclusive (for example, all of the traffic passing through an interface) or classification can be very specific (for example, you can use a class map with match commands that recognize specific aspects of the traffic). You can classify and apply QoS (for example, marking) and then, on another interface or network device, classify again based on the marked value and apply other QoS.

Policing

Policing determines whether a packet is in or out of profile by comparing the internal DSCP to the configured policer. Policer limits the bandwidth consumed by a traffic flow with the results given to the marker.

Policing and policers have the following attributes:

- Policers can occur only on a physical port basis.
- Policing can occur on ingress interfaces.
- Only one policer can be applied to a packet per direction.

Marking

Marking determines how to handle a packet when it is out of profile. It assesses the policer and the configuration data to determine the action required for the packet, and then handles the packet using one of the following methods:

- Let the packet through without modification
- Drop the packet

Marking can occur on ingress and egress interfaces.

Class Map

A class map names and isolates specific traffic from other traffic. The class map defines the criteria used to match against a specific traffic flow to classify it further. The criteria can include:

- Matching the access group defined by the ACL
- Matching a specific list of DSCP values

If there is more than one type of traffic to be classified, another class map can be created under a different name. After a packet is matched against the class-map criteria, it is further classified using a policy map.

Policy Map

A policy map specifies on which traffic class to act. This can be implemented as follows:

- Set a specific CoS or DSCP value in the traffic class.
- Specify the traffic bandwidth limitations for each matched traffic class (policer) and the action to take (marking) when the traffic is out of profile.

Policy maps have the following attributes:

- A policy map can contain multiple class statements, each with different match criteria and policers.
- A separate policy-map class can exist for each type of traffic received through an interface.
- There can be only one policy map per interface per direction. The same policy map can be applied to multiple interfaces and directions.
- Before a policy map can be effective, it must be attached to an interface.

MPLS Class

MPLS class or class specifies the class of the frames, for example frames with DSCP 0-7 belongs to class 0, DSCP 8-15 belongs to Class 1, and so on.

In OcNOS, there are 8 classes varying from 0-7. By default, EXP to class is mapped one-to-one.

For more information, refer the table below.

For MPLS Diff-Serv to work, QoS must be enabled at the global level. By default QoS is disabled.

Table 53. EXP to class mapping

| CoS | DSCP | EXP | Class | Queue |
|-----|-------|-----|-------|-------|
| 0 | 0-7 | 0 | 0 | 0 |
| 1 | 8-15 | 1 | 1 | 1 |
| 2 | 16-23 | 2 | 2 | 2 |
| 3 | 24-31 | 3 | 3 | 3 |
| 4 | 32-39 | 4 | 4 | 4 |
| 5 | 40-47 | 5 | 5 | 5 |
| 6 | 48-55 | 6 | 6 | 6 |
| 7 | 56-63 | 7 | 7 | 7 |

Remarking Configuration

This chapter contains a complete sample of configuring marking on global level and remarking of EXP bits on interface and global level along with LDP LSP for Pipe model and Uniform model.


Configuration

Configuring Remarking for MPLS EXP bits require the following configurations:

- Enabling label-switching on the interface on NSM.
- Configuring LSP (Using LDP, Static or RSVP-TE, in this example we are using LDP for setting up LSP).
- Running an IGP (Internal Gateway Protocol), for example, OSPF , to distribute reachability information within the MPLS cloud.
- Enable QoS, Configuring Remarking on interface and Global Level.

Topology

Figure 40. Basic LDP Topology



OSPF and LDP Configuration for PE1, P and PE2

PE1: NSM

| | |
|--------------------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#ip address 10.10.10.1/24 | Configure IP address for the interface |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |

| | |
|---|--|
| (config-if)#ip address 192.168.0.1/32 secondary | Set the IP address of the loopback interface to 192.168.0.1/32 |
| (config-if)#commit | Commit the transaction. |

PE1: OSPF

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.10.10.0/24 area 0 (config-router)#network 192.168.0.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

PE1: LDP

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 192.168.0.1 | Set the router ID to IP address 192.168.0.1. |
| (config-router)#transport-address ipv4 192.168.0.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe1. |
| (config-if)#commit | Commit the transaction. |

P: NSM

| | |
|--------------------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#ip address 10.10.10.2/24 | Configure IP address for the interface |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.11.1/24 | Configure IP address for the interface |

| | |
|---|--|
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 192.168.0.2/32 secondary | Set the IP address of the loopback interface to 192.168.0.2/32 |
| (config-if)#commit | Commit the transaction. |

P: OSPF

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.10.10.0/24 area 0 (config-router)#network 10.10.11.0/24 area 0 (config-router)#network 192.168.0.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

P: LDP

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 192.168.0.2 | Set the router ID to IP address 192.168.0.2. |
| (config-router)#transport-address ipv4 192.168.0.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| (config-if)#commit | Commit the transaction. |

PE2: NSM

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface xe1 | Specify the interface (xe2) to be configured. |
| (config-if)#ip address 10.10.11.2/24 | Configure IP address for the interface |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 192.168.0.3/32 secondary | Set the IP address of the loopback interface to 192.168.0.3/32 |
| (config-if)#commit | Commit the transaction. |

PE2: OSPF

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.10.11.0/24 area 0 (config-router)#network 192.168.0.3/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

PE2: LDP

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 192.168.0.3 | Set the router ID to IP address 192.168.0.3. |
| (config-router)#explicit-null | To disable PHP. |
| (config-router)#transport-address ipv4 192.168.0.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| (config-if)#commit | Commit the transaction. |

Configuration of Marking or Remarking

Marking can be done in Global level and Remarking can be done in Global level and in Interface level. Both methods are shown in the following sample configurations.

Global level configuration for PE1

| | |
|--|----------------------------|
| #configure terminal | Enter configure mode. |
| (config)#qos enable | Enable QOS. |
| (config)#qos profile exp-encap default | Configure default profile. |
| (config-egress-exp-encap-map)#l3 dscp 10 exp 3 | |
| (config-egress-exp-encap-map)#commit | |

Validation

```

PE1#show qos-profile type exp-encap default
profile name: default
profile type: exp-encap
configured mapping:
qos profile exp-encap default
  l3 dscp 10 exp 3
!
Detailed mapping:
  L3 DSCP to EXP
  -----+----- | -----+----- | -----+----- | -----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+----- | -----+-----
DSCP | EXP | DSCP | EXP | DSCP | EXP | DSCP | EXP
-----+----- | -----+----- | -----+----- | -----+-----
0 0 | 16 2 | 32 4 | 48 6
1 0 | 17 2 | 33 4 | 49 6
2 0 | 18 2 | 34 4 | 50 6
3 0 | 19 2 | 35 4 | 51 6
4 0 | 20 2 | 36 4 | 52 6
5 0 | 21 2 | 37 4 | 53 6
6 0 | 22 2 | 38 4 | 54 6
7 0 | 23 2 | 39 4 | 55 6
8 1 | 24 3 | 40 5 | 56 7
9 1 | 25 3 | 41 5 | 57 7
10 3 | 26 3 | 42 5 | 58 7
11 1 | 27 3 | 43 5 | 59 7
12 1 | 28 3 | 44 5 | 60 7
13 1 | 29 3 | 45 5 | 61 7
14 1 | 30 3 | 46 5 | 62 7
15 1 | 31 3 | 47 5 | 63 7

  L2 Queue + Color to EXP
  -----+----- | -----+----- | -----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+-----
Queue | Color | EXP | Queue | Color | EXP | Queue | Color | EXP
-----+-----+----- | -----+-----+----- | -----+-----+-----
0 green 0 | 0 yellow 0 | 0 red 0
1 green 1 | 1 yellow 1 | 1 red 1
2 green 2 | 2 yellow 2 | 2 red 2

```

```

3 green 3 | 3 yellow 3 | 3 red 3
4 green 4 | 4 yellow 4 | 4 red 4
5 green 5 | 5 yellow 5 | 5 red 5
6 green 6 | 6 yellow 6 | 6 red 6
7 green 7 | 7 yellow 7 | 7 red 7

```

Interface level configuration for PE1

| | |
|---|----------------------------------|
| #configure terminal | Enter configure mode. |
| (config)#qos enable | Enable QOS. |
| (config)#qos profile exp-encap TEST1 | Configure non-default profile. |
| (config-egress-exp-encap-map)#13 dscp 5 exp 5 | Map dscp value 5 to exp 5. |
| (config-egress-exp-encap-map)#exit | Exit exp-encap profile mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#qos map-profile exp-encap TEST1 | Mapping qos profile to interface |
| (config-if)#commit | Commit the transaction. |

Validation

```

PE1#show qos-profile type exp-encap TEST1
profile name: TEST1
profile type: exp-encap
profile attached to 1 instances
configured mapping:
 13 dscp 5 exp 5
Detailed mapping:
 L3 DSCP to EXP
-----+----- | -----+----- | -----+----- | -----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+----- | -----+-----
DSCP | EXP | DSCP | EXP | DSCP | EXP | DSCP | EXP
-----+----- | -----+----- | -----+----- | -----+-----
0 0 | 16 2 | 32 4 | 48 6
1 0 | 17 2 | 33 4 | 49 6
2 0 | 18 2 | 34 4 | 50 6
3 0 | 19 2 | 35 4 | 51 6
4 0 | 20 2 | 36 4 | 52 6
5 5 | 21 2 | 37 4 | 53 6
6 0 | 22 2 | 38 4 | 54 6
7 0 | 23 2 | 39 4 | 55 6
8 1 | 24 3 | 40 5 | 56 7
9 1 | 25 3 | 41 5 | 57 7
10 1 | 26 3 | 42 5 | 58 7
11 1 | 27 3 | 43 5 | 59 7
12 1 | 28 3 | 44 5 | 60 7
13 1 | 29 3 | 45 5 | 61 7
14 1 | 30 3 | 46 5 | 62 7
15 1 | 31 3 | 47 5 | 63 7

L2 Queue + Color to EXP

```

```
-----+----- | -----+----- | -----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+-----
Queue | Color | EXP | Queue | Color | EXP | Queue | Color | EXP
-----+-----+----- | -----+-----+----- | -----+-----+-----
0 green 0 | 0 yellow 0 | 0 red 0
1 green 1 | 1 yellow 1 | 1 red 1
2 green 2 | 2 yellow 2 | 2 red 2
3 green 3 | 3 yellow 3 | 3 red 3
4 green 4 | 4 yellow 4 | 4 red 4
5 green 5 | 5 yellow 5 | 5 red 5
6 green 6 | 6 yellow 6 | 6 red 6
7 green 7 | 7 yellow 7 | 7 red 7
```

Global level configuration for PE2

| | |
|--|-----------------------------|
| #configure terminal | Enter configure mode. |
| (config)#qos enable | Enable QOS. |
| (config)#qos statistics | Enable QOS statistics. |
| (config)#qos profile exp-to-queue default | Configure default profile |
| (config-ingress-exp-queue-map)#exp 5 queue 7 | Map exp value 5 to Class 7. |
| (config-ingress-exp-queue-map)#commit | Commit the transaction. |

Validation

PE2

```
PE2#show qos-profile type exp-to-queue

profile name: default
profile type: exp-to-queue
configured mapping:
qos profile exp-to-queue default
  exp 5 queue 7
!
Detailed mapping:
-----+-----
INPUT | OUTPUT
-----+-----
EXP | Queue | Color
-----+-----+-----
0 0 green
1 1 green
2 2 green
3 3 green
4 4 green
5 7 green
6 6 green
7 7 green
```

PE1

```

PE1#show qos-profile interface xe2
profile name: default
profile type: dscp-to-queue (Ingress)
mapping:
-----+----- | -----+----- | -----+-----
----- | -----+---
-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT | INPUT |
OUTPUT
-----+----- | -----+----- | -----+-----
----- | -----+---
-----
DSCP | TC | Color | Out DSCP | DSCP | TC | Color | Out DSCP | DSCP | TC | Color | Out DSCP | DSCP |
TC | Color | Out DSCP
-----+-----+-----+----- | -----+-----+-----+----- | -----+-----+-----+---
----- | -----+---
----+-----+-----
0 0 green 0 | 16 2 green 16 | 32 4 green 32 | 48 6
green 48
1 0 green 1 | 17 2 green 17 | 33 4 green 33 | 49 6
green 49
2 0 green 2 | 18 2 green 18 | 34 4 green 34 | 50 6
green 50
3 0 green 3 | 19 2 green 19 | 35 4 green 35 | 51 6
green 51
4 0 green 4 | 20 2 yellow 20 | 36 4 yellow 36 | 52 6
green 52
5 0 green 5 | 21 2 green 21 | 37 4 green 37 | 53 6
green 53
6 0 green 6 | 22 2 yellow 22 | 38 4 yellow 38 | 54 6
green 54
7 0 green 7 | 23 2 green 23 | 39 4 green 39 | 55 6
green 55
8 1 green 8 | 24 3 green 24 | 40 5 green 40 | 56 7
green 56
9 1 green 9 | 25 3 green 25 | 41 5 green 41 | 57 7
green 57
10 1 green 10 | 26 3 green 26 | 42 5 green 42 | 58 7
green 58
11 1 green 11 | 27 3 green 27 | 43 5 green 43 | 59 7
green 59
12 1 yellow 12 | 28 3 yellow 28 | 44 5 green 44 | 60 7
green 60
13 1 green 13 | 29 3 green 29 | 45 5 green 45 | 61 7
green 61
14 1 yellow 14 | 30 3 yellow 30 | 46 5 green 46 | 62 7
green 62
15 1 green 15 | 31 3 green 31 | 47 5 green 47 | 63 7
green 63

profile name: default
profile type: dscp-to-dscp (Egress)

```

Status: Inactive

mapping:

| INPUT | | | OUTPUT | | | INPUT | | | OUTPUT | | |
|-------|-------|----------|--------|--------|----------|-------|-------|----------|--------|-------|----------|
| DSCP | Color | Out DSCP | DSCP | Color | Out DSCP | DSCP | Color | Out DSCP | DSCP | Color | Out DSCP |
| 0 | green | 0 | 0 | yellow | 0 | 0 | red | 0 | | | |
| 1 | green | 1 | 1 | yellow | 1 | 1 | red | 1 | | | |
| 2 | green | 2 | 2 | yellow | 2 | 2 | red | 2 | | | |
| 3 | green | 3 | 3 | yellow | 3 | 3 | red | 3 | | | |
| 4 | green | 4 | 4 | yellow | 4 | 4 | red | 4 | | | |
| 5 | green | 5 | 5 | yellow | 5 | 5 | red | 5 | | | |
| 6 | green | 6 | 6 | yellow | 6 | 6 | red | 6 | | | |
| 7 | green | 7 | 7 | yellow | 7 | 7 | red | 7 | | | |
| 8 | green | 8 | 8 | yellow | 8 | 8 | red | 8 | | | |
| 9 | green | 9 | 9 | yellow | 9 | 9 | red | 9 | | | |
| 10 | green | 10 | 10 | yellow | 12 | 10 | red | 14 | | | |
| 11 | green | 11 | 11 | yellow | 11 | 11 | red | 11 | | | |
| 12 | green | 12 | 12 | yellow | 12 | 12 | red | 14 | | | |
| 13 | green | 13 | 13 | yellow | 13 | 13 | red | 13 | | | |
| 14 | green | 14 | 14 | yellow | 14 | 14 | red | 14 | | | |
| 15 | green | 15 | 15 | yellow | 15 | 15 | red | 15 | | | |
| 16 | green | 16 | 16 | yellow | 16 | 16 | red | 16 | | | |
| 17 | green | 17 | 17 | yellow | 17 | 17 | red | 17 | | | |
| 18 | green | 18 | 18 | yellow | 20 | 18 | red | 22 | | | |
| 19 | green | 19 | 19 | yellow | 19 | 19 | red | 19 | | | |
| 20 | green | 20 | 20 | yellow | 20 | 20 | red | 22 | | | |
| 21 | green | 21 | 21 | yellow | 21 | 21 | red | 21 | | | |
| 22 | green | 22 | 22 | yellow | 22 | 22 | red | 22 | | | |
| 23 | green | 23 | 23 | yellow | 23 | 23 | red | 23 | | | |
| 24 | green | 24 | 24 | yellow | 24 | 24 | red | 24 | | | |
| 25 | green | 25 | 25 | yellow | 25 | 25 | red | 25 | | | |
| 26 | green | 26 | 26 | yellow | 28 | 26 | red | 30 | | | |
| 27 | green | 27 | 27 | yellow | 27 | 27 | red | 27 | | | |
| 28 | green | 28 | 28 | yellow | 28 | 28 | red | 30 | | | |
| 29 | green | 29 | 29 | yellow | 29 | 29 | red | 29 | | | |
| 30 | green | 30 | 30 | yellow | 30 | 30 | red | 30 | | | |
| 31 | green | 31 | 31 | yellow | 31 | 31 | red | 31 | | | |
| 32 | green | 32 | 32 | yellow | 32 | 32 | red | 32 | | | |
| 33 | green | 33 | 33 | yellow | 33 | 33 | red | 33 | | | |
| 34 | green | 34 | 34 | yellow | 36 | 34 | red | 38 | | | |
| 35 | green | 35 | 35 | yellow | 35 | 35 | red | 35 | | | |
| 36 | green | 36 | 36 | yellow | 36 | 36 | red | 38 | | | |
| 37 | green | 37 | 37 | yellow | 37 | 37 | red | 37 | | | |
| 38 | green | 38 | 38 | yellow | 38 | 38 | red | 38 | | | |
| 39 | green | 39 | 39 | yellow | 39 | 39 | red | 39 | | | |
| 40 | green | 40 | 40 | yellow | 40 | 40 | red | 40 | | | |
| 41 | green | 41 | 41 | yellow | 41 | 41 | red | 41 | | | |
| 42 | green | 42 | 42 | yellow | 42 | 42 | red | 42 | | | |
| 43 | green | 43 | 43 | yellow | 43 | 43 | red | 43 | | | |
| 44 | green | 44 | 44 | yellow | 44 | 44 | red | 44 | | | |
| 45 | green | 45 | 45 | yellow | 45 | 45 | red | 45 | | | |
| 46 | green | 46 | 46 | yellow | 46 | 46 | red | 46 | | | |
| 47 | green | 47 | 47 | yellow | 47 | 47 | red | 47 | | | |
| 48 | green | 48 | 48 | yellow | 48 | 48 | red | 48 | | | |
| 49 | green | 49 | 49 | yellow | 49 | 49 | red | 49 | | | |
| 50 | green | 50 | 50 | yellow | 50 | 50 | red | 50 | | | |
| 51 | green | 51 | 51 | yellow | 51 | 51 | red | 51 | | | |
| 52 | green | 52 | 52 | yellow | 52 | 52 | red | 52 | | | |
| 53 | green | 53 | 53 | yellow | 53 | 53 | red | 53 | | | |
| 54 | green | 54 | 54 | yellow | 54 | 54 | red | 54 | | | |
| 55 | green | 55 | 55 | yellow | 55 | 55 | red | 55 | | | |
| 56 | green | 56 | 56 | yellow | 56 | 56 | red | 56 | | | |
| 57 | green | 57 | 57 | yellow | 57 | 57 | red | 57 | | | |

```

58 green 58 | 58 yellow 58 | 58 red 58
59 green 59 | 59 yellow 59 | 59 red 59
60 green 60 | 60 yellow 60 | 60 red 60
61 green 61 | 61 yellow 61 | 61 red 61
62 green 62 | 62 yellow 62 | 62 red 62
63 green 63 | 63 yellow 63 | 63 red 63

```

```

profile name: default
profile type: exp-to-queue (Ingress)
mapping:

```

```

-----+-----
INPUT | OUTPUT
-----+-----

```

```

EXP | TC | Color
-----+-----

```

```

0 0 green
1 1 green
2 2 green
3 3 green
4 4 green
5 5 green
6 6 green
7 7 green

```

```

profile name: TEST1
profile type: exp-encap (Egress)
mapping:

```

```

L3 DSCP to EXP

```

```

-----+----- | -----+----- | -----+----- | -----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+----- | -----+-----
DSCP | EXP | DSCP | EXP | DSCP | EXP | DSCP | EXP
-----+----- | -----+----- | -----+----- | -----+-----
0 0 | 16 2 | 32 4 | 48 6
1 0 | 17 2 | 33 4 | 49 6
2 0 | 18 2 | 34 4 | 50 6
3 0 | 19 2 | 35 4 | 51 6
4 0 | 20 2 | 36 4 | 52 6
5 5 | 21 2 | 37 4 | 53 6
6 0 | 22 2 | 38 4 | 54 6
7 0 | 23 2 | 39 4 | 55 6
8 1 | 24 3 | 40 5 | 56 7
9 1 | 25 3 | 41 5 | 57 7
10 1 | 26 3 | 42 5 | 58 7
11 1 | 27 3 | 43 5 | 59 7
12 1 | 28 3 | 44 5 | 60 7
13 1 | 29 3 | 45 5 | 61 7
14 1 | 30 3 | 46 5 | 62 7
15 1 | 31 3 | 47 5 | 63 7

```

```

L2 Traffic-Class + Color to EXP

```

```

-----+----- | -----+----- | -----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+-----
TC | Color | EXP | TC | Color | EXP | TC | Color | EXP
-----+----- | -----+----- | -----+-----
0 green 0 | 0 yellow 0 | 0 red 0
1 green 1 | 1 yellow 1 | 1 red 1
2 green 2 | 2 yellow 2 | 2 red 2
3 green 3 | 3 yellow 3 | 3 red 3
4 green 4 | 4 yellow 4 | 4 red 4
5 green 5 | 5 yellow 5 | 5 red 5
6 green 6 | 6 yellow 6 | 6 red 6
7 green 7 | 7 yellow 7 | 7 red 7

```

```

PE1#show qos TEST1
profile name: TEST1
profile type: exp-encap
profile attached to 1 instances
configured mapping:
  l3 dscp 5 exp 5
Detailed mapping:
  L3 DSCP to EXP
  -----+----- | -----+----- | -----+----- | -----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+----- | -----+-----
DSCP | EXP | DSCP | EXP | DSCP | EXP | DSCP | EXP
-----+----- | -----+----- | -----+----- | -----+-----
0 0 | 16 2 | 32 4 | 48 6
1 0 | 17 2 | 33 4 | 49 6
2 0 | 18 2 | 34 4 | 50 6
3 0 | 19 2 | 35 4 | 51 6
4 0 | 20 2 | 36 4 | 52 6
5 5 | 21 2 | 37 4 | 53 6
6 0 | 22 2 | 38 4 | 54 6
7 0 | 23 2 | 39 4 | 55 6
8 1 | 24 3 | 40 5 | 56 7
9 1 | 25 3 | 41 5 | 57 7
10 1 | 26 3 | 42 5 | 58 7
11 1 | 27 3 | 43 5 | 59 7
12 1 | 28 3 | 44 5 | 60 7
13 1 | 29 3 | 45 5 | 61 7
14 1 | 30 3 | 46 5 | 62 7
15 1 | 31 3 | 47 5 | 63 7

  L2 Traffic-Class + Color to EXP
  -----+----- | -----+----- | -----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+-----
TC | Color | EXP | TC | Color | EXP | TC | Color | EXP
-----+----- | -----+----- | -----+-----
0 green 0 | 0 yellow 0 | 0 red 0
1 green 1 | 1 yellow 1 | 1 red 1
2 green 2 | 2 yellow 2 | 2 red 2
3 green 3 | 3 yellow 3 | 3 red 3
4 green 4 | 4 yellow 4 | 4 red 4
5 green 5 | 5 yellow 5 | 5 red 5
6 green 6 | 6 yellow 6 | 6 red 6
7 green 7 | 7 yellow 7 | 7 red 7

```

PE2

```

PE2#show qos int xel
profile name: default
profile type: dscp-to-queue (Ingress)
mapping:
  -----+----- | -----+----- | -----+-----
  ----- | -----+-----
  -----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT | INPUT |
OUTPUT
  -----+----- | -----+----- | -----+-----
  ----- | -----+-----
  -----
DSCP | TC | Color | Out DSCP | DSCP | TC | Color | Out DSCP | DSCP | TC | Color | Out DSCP | DSCP |
TC | Color | Out DSCP

```

```

-----+-----+-----+----- | -----+-----+-----+----- | -----+-----+-----+-----
----- | -----+-----
-----+-----+-----
0 0 green 0 | 16 2 green 16 | 32 4 green 32 | 48 6
green 48
1 0 green 1 | 17 2 green 17 | 33 4 green 33 | 49 6
green 49
2 0 green 2 | 18 2 green 18 | 34 4 green 34 | 50 6
green 50
3 0 green 3 | 19 2 green 19 | 35 4 green 35 | 51 6
green 51
4 0 green 4 | 20 2 yellow 20 | 36 4 yellow 36 | 52 6
green 52
5 0 green 5 | 21 2 green 21 | 37 4 green 37 | 53 6
green 53
6 0 green 6 | 22 2 yellow 22 | 38 4 yellow 38 | 54 6
green 54
7 0 green 7 | 23 2 green 23 | 39 4 green 39 | 55 6
green 55
8 1 green 8 | 24 3 green 24 | 40 5 green 40 | 56 7
green 56
9 1 green 9 | 25 3 green 25 | 41 5 green 41 | 57 7
green 57
10 1 green 10 | 26 3 green 26 | 42 5 green 42 | 58 7
green 58
11 1 green 11 | 27 3 green 27 | 43 5 green 43 | 59 7
green 59
12 1 yellow 12 | 28 3 yellow 28 | 44 5 green 44 | 60 7
green 60
13 1 green 13 | 29 3 green 29 | 45 5 green 45 | 61 7
green 61
14 1 yellow 14 | 30 3 yellow 30 | 46 5 green 46 | 62 7
green 62
15 1 green 15 | 31 3 green 31 | 47 5 green 47 | 63 7
green 63

profile name: default
profile type: dscp-to-dscp (Egress)
Status: Inactive
mapping:
-----+----- | -----+----- | -----+-----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+-----+-----
DSCP | Color | Out DSCP | DSCP | Color | Out DSCP | DSCP | Color | Out DSCP
-----+-----+----- | -----+-----+----- | -----+-----+-----
0 green 0 | 0 yellow 0 | 0 red 0
1 green 1 | 1 yellow 1 | 1 red 1
2 green 2 | 2 yellow 2 | 2 red 2
3 green 3 | 3 yellow 3 | 3 red 3
4 green 4 | 4 yellow 4 | 4 red 4
5 green 5 | 5 yellow 5 | 5 red 5
6 green 6 | 6 yellow 6 | 6 red 6
7 green 7 | 7 yellow 7 | 7 red 7
8 green 8 | 8 yellow 8 | 8 red 8
9 green 9 | 9 yellow 9 | 9 red 9
10 green 10 | 10 yellow 12 | 10 red 14
11 green 11 | 11 yellow 11 | 11 red 11
12 green 12 | 12 yellow 12 | 12 red 14
13 green 13 | 13 yellow 13 | 13 red 13
14 green 14 | 14 yellow 14 | 14 red 14
15 green 15 | 15 yellow 15 | 15 red 15
16 green 16 | 16 yellow 16 | 16 red 16
17 green 17 | 17 yellow 17 | 17 red 17
18 green 18 | 18 yellow 20 | 18 red 22
19 green 19 | 19 yellow 19 | 19 red 19

```



```

20 green 20 | 20 yellow 20 | 20 red 22
21 green 21 | 21 yellow 21 | 21 red 21
22 green 22 | 22 yellow 22 | 22 red 22
23 green 23 | 23 yellow 23 | 23 red 23
24 green 24 | 24 yellow 24 | 24 red 24
25 green 25 | 25 yellow 25 | 25 red 25
26 green 26 | 26 yellow 28 | 26 red 30
27 green 27 | 27 yellow 27 | 27 red 27
28 green 28 | 28 yellow 28 | 28 red 30
29 green 29 | 29 yellow 29 | 29 red 29
30 green 30 | 30 yellow 30 | 30 red 30
31 green 31 | 31 yellow 31 | 31 red 31
32 green 32 | 32 yellow 32 | 32 red 32
33 green 33 | 33 yellow 33 | 33 red 33
34 green 34 | 34 yellow 36 | 34 red 38
35 green 35 | 35 yellow 35 | 35 red 35
36 green 36 | 36 yellow 36 | 36 red 38
37 green 37 | 37 yellow 37 | 37 red 37
38 green 38 | 38 yellow 38 | 38 red 38
39 green 39 | 39 yellow 39 | 39 red 39
40 green 40 | 40 yellow 40 | 40 red 40
41 green 41 | 41 yellow 41 | 41 red 41
42 green 42 | 42 yellow 42 | 42 red 42
43 green 43 | 43 yellow 43 | 43 red 43
44 green 44 | 44 yellow 44 | 44 red 44
45 green 45 | 45 yellow 45 | 45 red 45
46 green 46 | 46 yellow 46 | 46 red 46
47 green 47 | 47 yellow 47 | 47 red 47
48 green 48 | 48 yellow 48 | 48 red 48
49 green 49 | 49 yellow 49 | 49 red 49
50 green 50 | 50 yellow 50 | 50 red 50
51 green 51 | 51 yellow 51 | 51 red 51
52 green 52 | 52 yellow 52 | 52 red 52
53 green 53 | 53 yellow 53 | 53 red 53
54 green 54 | 54 yellow 54 | 54 red 54
55 green 55 | 55 yellow 55 | 55 red 55
56 green 56 | 56 yellow 56 | 56 red 56
57 green 57 | 57 yellow 57 | 57 red 57
58 green 58 | 58 yellow 58 | 58 red 58
59 green 59 | 59 yellow 59 | 59 red 59
60 green 60 | 60 yellow 60 | 60 red 60
61 green 61 | 61 yellow 61 | 61 red 61
62 green 62 | 62 yellow 62 | 62 red 62
63 green 63 | 63 yellow 63 | 63 red 63

```

```

profile name: default
profile type: exp-to-queue (Ingress)
mapping:
-----+-----
INPUT | OUTPUT
-----+-----
EXP | TC | Color
-----+-----+-----
0 0 green
1 1 green
2 2 green
3 3 green
4 4 green
5 7 green
6 6 green
7 7 green

```

```

profile name: default
profile type: exp-encap (Egress)
mapping:

```

L3 DSCP to EXP

| INPUT | | OUTPUT | | INPUT | | OUTPUT | | INPUT | | OUTPUT | |
|-------|------|--------|------|-------|-----|--------|-----|-------|-----|--------|-----|
| DSCP | EXP | DSCP | EXP | DSCP | EXP | DSCP | EXP | DSCP | EXP | DSCP | EXP |
| 0 0 | 16 2 | 32 4 | 48 6 | | | | | | | | |
| 1 0 | 17 2 | 33 4 | 49 6 | | | | | | | | |
| 2 0 | 18 2 | 34 4 | 50 6 | | | | | | | | |
| 3 0 | 19 2 | 35 4 | 51 6 | | | | | | | | |
| 4 0 | 20 2 | 36 4 | 52 6 | | | | | | | | |
| 5 0 | 21 2 | 37 4 | 53 6 | | | | | | | | |
| 6 0 | 22 2 | 38 4 | 54 6 | | | | | | | | |
| 7 0 | 23 2 | 39 4 | 55 6 | | | | | | | | |
| 8 1 | 24 3 | 40 5 | 56 7 | | | | | | | | |
| 9 1 | 25 3 | 41 5 | 57 7 | | | | | | | | |
| 10 1 | 26 3 | 42 5 | 58 7 | | | | | | | | |
| 11 1 | 27 3 | 43 5 | 59 7 | | | | | | | | |
| 12 1 | 28 3 | 44 5 | 60 7 | | | | | | | | |
| 13 1 | 29 3 | 45 5 | 61 7 | | | | | | | | |
| 14 1 | 30 3 | 46 5 | 62 7 | | | | | | | | |
| 15 1 | 31 3 | 47 5 | 63 7 | | | | | | | | |

L2 Traffic-Class + Color to EXP

| INPUT | | OUTPUT | | INPUT | | OUTPUT | |
|-----------|------------|---------|----|-------|-----|--------|-------|
| TC | Color | EXP | TC | Color | EXP | TC | Color |
| 0 green 0 | 0 yellow 0 | 0 red 0 | | | | | |
| 1 green 1 | 1 yellow 1 | 1 red 1 | | | | | |
| 2 green 2 | 2 yellow 2 | 2 red 2 | | | | | |
| 3 green 3 | 3 yellow 3 | 3 red 3 | | | | | |
| 4 green 4 | 4 yellow 4 | 4 red 4 | | | | | |
| 5 green 5 | 5 yellow 5 | 5 red 5 | | | | | |
| 6 green 6 | 6 yellow 6 | 6 red 6 | | | | | |
| 7 green 7 | 7 yellow 7 | 7 red 7 | | | | | |

```
PE2#show qos default
```

```
profile name: default
```

```
profile type: cos-to-queue
```

```
configured mapping:
```

```
Detailed mapping:
```

| INPUT | | OUTPUT | | INPUT | | OUTPUT | |
|-------------|--------------|--------|-------|-------|-----|--------|-------|
| COS | DEI | TC | Color | COS | DEI | TC | Color |
| 0 0 0 green | 0 1 0 yellow | | | | | | |
| 1 0 1 green | 1 1 1 yellow | | | | | | |
| 2 0 2 green | 2 1 2 yellow | | | | | | |
| 3 0 3 green | 3 1 3 yellow | | | | | | |
| 4 0 4 green | 4 1 4 yellow | | | | | | |
| 5 0 5 green | 5 1 5 yellow | | | | | | |
| 6 0 6 green | 6 1 6 yellow | | | | | | |
| 7 0 7 green | 7 1 7 yellow | | | | | | |

```
profile name: default
```

```
profile type: queue-color-to-cos
```

```
configured mapping:
```

```
Detailed mapping:
```

| INPUT | | OUTPUT | | INPUT | | OUTPUT | |
|-------|-------|--------|----|-------|-----|--------|-------|
| TC | Color | COS | TC | Color | COS | TC | Color |

```

-----+-----+----- | -----+-----+----- | -----+-----+-----
0 green 0 | 0 yellow 0 | 0 red 0
1 green 1 | 1 yellow 1 | 1 red 1
2 green 2 | 2 yellow 2 | 2 red 2
3 green 3 | 3 yellow 3 | 3 red 3
4 green 4 | 4 yellow 4 | 4 red 4
5 green 5 | 5 yellow 5 | 5 red 5
6 green 6 | 6 yellow 6 | 6 red 6
7 green 7 | 7 yellow 7 | 7 red 7

profile name: default
profile type: dscp-to-queue
configured mapping:
Detailed mapping:
-----+-----+----- | -----+-----+----- | -----+-----+-----
----- | -----+---
-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT | INPUT |
OUTPUT
-----+-----+----- | -----+-----+----- | -----+-----+-----
----- | -----+---
-----
DSCP | TC | Color | Out DSCP | DSCP | TC | Color | Out DSCP | DSCP | TC | Color | Out DSCP | DSCP |
TC | Color | Out DSCP
-----+-----+-----+----- | -----+-----+-----+----- | -----+-----+-----+-----
----- | -----+---
----+-----+-----
0 0 green 0 | 16 2 green 16 | 32 4 green 32 | 48 6
green 48
1 0 green 1 | 17 2 green 17 | 33 4 green 33 | 49 6
green 49
2 0 green 2 | 18 2 green 18 | 34 4 green 34 | 50 6
green 50
3 0 green 3 | 19 2 green 19 | 35 4 green 35 | 51 6
green 51
4 0 green 4 | 20 2 yellow 20 | 36 4 yellow 36 | 52 6
green 52
5 0 green 5 | 21 2 green 21 | 37 4 green 37 | 53 6
green 53
6 0 green 6 | 22 2 yellow 22 | 38 4 yellow 38 | 54 6
green 54
7 0 green 7 | 23 2 green 23 | 39 4 green 39 | 55 6
green 55
8 1 green 8 | 24 3 green 24 | 40 5 green 40 | 56 7
green 56
9 1 green 9 | 25 3 green 25 | 41 5 green 41 | 57 7
green 57
10 1 green 10 | 26 3 green 26 | 42 5 green 42 | 58 7
green 58
11 1 green 11 | 27 3 green 27 | 43 5 green 43 | 59 7
green 59
12 1 yellow 12 | 28 3 yellow 28 | 44 5 green 44 | 60 7
green 60
13 1 green 13 | 29 3 green 29 | 45 5 green 45 | 61 7
green 61
14 1 yellow 14 | 30 3 yellow 30 | 46 5 green 46 | 62 7
green 62
15 1 green 15 | 31 3 green 31 | 47 5 green 47 | 63 7
green 63

profile name: default
profile type: queue-color-to-dscp
configured mapping:
Detailed mapping:
-----+-----+----- | -----+-----+----- | -----+-----+-----

```

```

INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+----- | -----+----- | -----+-----
TC | Color | DSCP | TC | Color | DSCP | TC | Color | DSCP
-----+-----+----- | -----+-----+----- | -----+-----+-----
0 green 0 | 0 yellow 0 | 0 red 0
1 green 10 | 1 yellow 12 | 1 red 14
2 green 18 | 2 yellow 20 | 2 red 22
3 green 26 | 3 yellow 28 | 3 red 30
4 green 34 | 4 yellow 36 | 4 red 38
5 green 40 | 5 yellow 40 | 5 red 40
6 green 48 | 6 yellow 48 | 6 red 48
7 green 56 | 7 yellow 56 | 7 red 56

profile name: default
profile type: dscp-to-dscp
configured mapping:
Detailed mapping:
-----+-----+----- | -----+-----+----- | -----+-----+-----
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
-----+-----+----- | -----+-----+----- | -----+-----+-----
DSCP | Color | Out DSCP | DSCP | Color | Out DSCP | DSCP | Color | Out DSCP
-----+-----+----- | -----+-----+----- | -----+-----+-----
0 green 0 | 0 yellow 0 | 0 red 0
1 green 1 | 1 yellow 1 | 1 red 1
2 green 2 | 2 yellow 2 | 2 red 2
3 green 3 | 3 yellow 3 | 3 red 3
4 green 4 | 4 yellow 4 | 4 red 4
5 green 5 | 5 yellow 5 | 5 red 5
6 green 6 | 6 yellow 6 | 6 red 6
7 green 7 | 7 yellow 7 | 7 red 7
8 green 8 | 8 yellow 8 | 8 red 8
9 green 9 | 9 yellow 9 | 9 red 9
10 green 10 | 10 yellow 12 | 10 red 14
11 green 11 | 11 yellow 11 | 11 red 11
12 green 12 | 12 yellow 12 | 12 red 14
13 green 13 | 13 yellow 13 | 13 red 13
14 green 14 | 14 yellow 14 | 14 red 14
15 green 15 | 15 yellow 15 | 15 red 15
16 green 16 | 16 yellow 16 | 16 red 16
17 green 17 | 17 yellow 17 | 17 red 17
18 green 18 | 18 yellow 20 | 18 red 22
19 green 19 | 19 yellow 19 | 19 red 19
20 green 20 | 20 yellow 20 | 20 red 22
21 green 21 | 21 yellow 21 | 21 red 21
22 green 22 | 22 yellow 22 | 22 red 22
23 green 23 | 23 yellow 23 | 23 red 23
24 green 24 | 24 yellow 24 | 24 red 24
25 green 25 | 25 yellow 25 | 25 red 25
26 green 26 | 26 yellow 28 | 26 red 30
27 green 27 | 27 yellow 27 | 27 red 27
28 green 28 | 28 yellow 28 | 28 red 30
29 green 29 | 29 yellow 29 | 29 red 29
30 green 30 | 30 yellow 30 | 30 red 30
31 green 31 | 31 yellow 31 | 31 red 31
32 green 32 | 32 yellow 32 | 32 red 32
33 green 33 | 33 yellow 33 | 33 red 33
34 green 34 | 34 yellow 36 | 34 red 38
35 green 35 | 35 yellow 35 | 35 red 35
36 green 36 | 36 yellow 36 | 36 red 38
37 green 37 | 37 yellow 37 | 37 red 37
38 green 38 | 38 yellow 38 | 38 red 38
39 green 39 | 39 yellow 39 | 39 red 39
40 green 40 | 40 yellow 40 | 40 red 40
41 green 41 | 41 yellow 41 | 41 red 41
42 green 42 | 42 yellow 42 | 42 red 42

```

```

43 green 43 | 43 yellow 43 | 43 red 43
44 green 44 | 44 yellow 44 | 44 red 44
45 green 45 | 45 yellow 45 | 45 red 45
46 green 46 | 46 yellow 46 | 46 red 46
47 green 47 | 47 yellow 47 | 47 red 47
48 green 48 | 48 yellow 48 | 48 red 48
49 green 49 | 49 yellow 49 | 49 red 49
50 green 50 | 50 yellow 50 | 50 red 50
51 green 51 | 51 yellow 51 | 51 red 51
52 green 52 | 52 yellow 52 | 52 red 52
53 green 53 | 53 yellow 53 | 53 red 53
54 green 54 | 54 yellow 54 | 54 red 54
55 green 55 | 55 yellow 55 | 55 red 55
56 green 56 | 56 yellow 56 | 56 red 56
57 green 57 | 57 yellow 57 | 57 red 57
58 green 58 | 58 yellow 58 | 58 red 58
59 green 59 | 59 yellow 59 | 59 red 59
60 green 60 | 60 yellow 60 | 60 red 60
61 green 61 | 61 yellow 61 | 61 red 61
62 green 62 | 62 yellow 62 | 62 red 62
63 green 63 | 63 yellow 63 | 63 red 63

```

```

profile name: default
profile type: dscp-encap
configured mapping:
Detailed mapping:

```

```
L3 DSCP to DSCP-ENCAP
```

| INPUT | | OUTPUT | | INPUT | | OUTPUT | | INPUT | | OUTPUT | | INPUT | | OUTPUT | |
|-------|-------|--------|-------|-------|------|--------|------|-------|------|--------|------|-------|------|--------|------|
| DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP | DSCP |
| 0 0 | 16 16 | 32 32 | 48 48 | | | | | | | | | | | | |
| 1 1 | 17 17 | 33 33 | 49 49 | | | | | | | | | | | | |
| 2 2 | 18 18 | 34 34 | 50 50 | | | | | | | | | | | | |
| 3 3 | 19 19 | 35 35 | 51 51 | | | | | | | | | | | | |
| 4 4 | 20 20 | 36 36 | 52 52 | | | | | | | | | | | | |
| 5 5 | 21 21 | 37 37 | 53 53 | | | | | | | | | | | | |
| 6 6 | 22 22 | 38 38 | 54 54 | | | | | | | | | | | | |
| 7 7 | 23 23 | 39 39 | 55 55 | | | | | | | | | | | | |
| 8 8 | 24 24 | 40 40 | 56 56 | | | | | | | | | | | | |
| 9 9 | 25 25 | 41 41 | 57 57 | | | | | | | | | | | | |
| 10 10 | 26 26 | 42 42 | 58 58 | | | | | | | | | | | | |
| 11 11 | 27 27 | 43 43 | 59 59 | | | | | | | | | | | | |
| 12 12 | 28 28 | 44 44 | 60 60 | | | | | | | | | | | | |
| 13 13 | 29 29 | 45 45 | 61 61 | | | | | | | | | | | | |
| 14 14 | 30 30 | 46 46 | 62 62 | | | | | | | | | | | | |
| 15 15 | 31 31 | 47 47 | 63 63 | | | | | | | | | | | | |

```
L2 Traffic-Class + Color to DSCP-ENCAP
```

| INPUT | | OUTPUT | | INPUT | | OUTPUT | | INPUT | | OUTPUT | |
|-------|----------|--------|-----------|-------|--------|--------|-------|-------|----|--------|------|
| TC | Color | DSCP | TC | Color | DSCP | TC | Color | DSCP | TC | Color | DSCP |
| 0 | green 0 | 0 | yellow 0 | 0 | red 0 | | | | | | |
| 1 | green 8 | 1 | yellow 8 | 1 | red 8 | | | | | | |
| 2 | green 16 | 2 | yellow 16 | 2 | red 16 | | | | | | |
| 3 | green 24 | 3 | yellow 24 | 3 | red 24 | | | | | | |
| 4 | green 32 | 4 | yellow 32 | 4 | red 32 | | | | | | |
| 5 | green 40 | 5 | yellow 40 | 5 | red 40 | | | | | | |
| 6 | green 48 | 6 | yellow 48 | 6 | red 48 | | | | | | |
| 7 | green 56 | 7 | yellow 56 | 7 | red 56 | | | | | | |

```
profile name: default
```

```

profile type: exp-to-queue
configured mapping:
qos profile exp-to-queue default
  exp 5 queue 7
!

```

Detailed mapping:

```
-----+-----
```

```
INPUT | OUTPUT
```

```
-----+-----
```

```
EXP | TC | Color
```

```
-----+-----+-----
```

```
0 0 green
```

```
1 1 green
```

```
2 2 green
```

```
3 3 green
```

```
4 4 green
```

```
5 7 green
```

```
6 6 green
```

```
7 7 green
```

profile name: default

profile type: exp-encap

configured mapping:

Detailed mapping:

L3 DSCP to EXP

```
-----+----- | -----+----- | -----+----- | -----+-----
```

```
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
```

```
-----+----- | -----+----- | -----+----- | -----+-----
```

```
DSCP | EXP | DSCP | EXP | DSCP | EXP | DSCP | EXP
```

```
-----+----- | -----+----- | -----+----- | -----+-----
```

```
0 0 | 16 2 | 32 4 | 48 6
```

```
1 0 | 17 2 | 33 4 | 49 6
```

```
2 0 | 18 2 | 34 4 | 50 6
```

```
3 0 | 19 2 | 35 4 | 51 6
```

```
4 0 | 20 2 | 36 4 | 52 6
```

```
5 0 | 21 2 | 37 4 | 53 6
```

```
6 0 | 22 2 | 38 4 | 54 6
```

```
7 0 | 23 2 | 39 4 | 55 6
```

```
8 1 | 24 3 | 40 5 | 56 7
```

```
9 1 | 25 3 | 41 5 | 57 7
```

```
10 1 | 26 3 | 42 5 | 58 7
```

```
11 1 | 27 3 | 43 5 | 59 7
```

```
12 1 | 28 3 | 44 5 | 60 7
```

```
13 1 | 29 3 | 45 5 | 61 7
```

```
14 1 | 30 3 | 46 5 | 62 7
```

```
15 1 | 31 3 | 47 5 | 63 7
```

L2 Traffic-Class + Color to EXP

```
-----+----- | -----+----- | -----+-----
```

```
INPUT | OUTPUT | INPUT | OUTPUT | INPUT | OUTPUT
```

```
-----+----- | -----+----- | -----+-----
```

```
TC | Color | EXP | TC | Color | EXP | TC | Color | EXP
```

```
-----+-----+----- | -----+-----+----- | -----+-----+-----
```

```
0 green 0 | 0 yellow 0 | 0 red 0
```

```
1 green 1 | 1 yellow 1 | 1 red 1
```

```
2 green 2 | 2 yellow 2 | 2 red 2
```

```
3 green 3 | 3 yellow 3 | 3 red 3
```

```
4 green 4 | 4 yellow 4 | 4 red 4
```

```
5 green 5 | 5 yellow 5 | 5 red 5
```

```
6 green 6 | 6 yellow 6 | 6 red 6
```

```
7 green 7 | 7 yellow 7 | 7 red 7
```

Policing Configuration

This chapter contains a complete sample of configuration of Policing for Pipe and Uniform models. This example shows configurations using LDP.


Configuration

Configuring Remarking for MPLS EXP bits require the following configurations:

- Enabling label-switching on the interface on NSM.
- Configuring LSP (Using LDP, Static or RSVP-TE, in this example we are using LDP for setting UP LSP).
- Running an IGP (Internal Gateway Protocol), for example, OSPF, to distribute reachability information within the MPLS cloud.
- Enable QoS, Configuring Policing on interface Level.

Topology

Figure 41. Basic Policing Topology



The following steps describe how to configure Policing.

PE1:NSM

| | |
|--------------------------------------|---|
| #configure terminal | Enter configure mode |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.10.1/24 | Configure IP address for the interface |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |

| | |
|--|--|
| (config-if)#ip address 192.168.0.1/32 secondary | Set the IP address of the loopback interface to 192.168.0.1/32 |
| (config-if)#commit | Commit the transaction. |

PE1:OSPF

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.10.10.0/24 area 0 (config-router)#network 192.168.0.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

PE1:LDP

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 192.168.0.1 | Set the router ID to IP address 192.168.0.1. |
| (config-router)#transport-address ipv4 192.168.0.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe1. |
| (config-if)#commit | Commit the transaction. |

PE1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#qos enable | Enable QoS. |
| (config)#hardware-profile filter qos enable | Enable QoS hardware. |
| (config)#commit | Commit candidate configuration to the running configuration. |
| (config)#class-map cmap1 | Enter Class-map mode. |
| (config-cmap-qos)#match dscp 2 | Configure match criteria as DSCP with Value 2. |

| | |
|---|--|
| (config-cmap-qos)#exit | Exit Class-map mode. |
| (config)#commit | Commit candidate configuration to the running configuration. |
| (config)#policy-map pmap1 | Enter policy-map mode. |
| (config-pmap-qos)#class cmap1 | Assign Class cmap1 to Policy-map pmap1. |
| (config-pmap-c-qos)#police cir 1 mbps eir 1 mbps bc 1 kbytes be 1 kbytes | Police DSCP 2 packets @ Committed information rate 1 mbps. |
| (config-pmap-c-qos)#exit | Exit out of policy-class-map mode. |
| (config-pmap-qos)#exit | Exit out of Policy-map mode. |
| (config)#commit | Commit candidate configuration to the running configuration. |
| (config)#interface xe1 | Enter xe1 interface. |
| (config-if)#service-policy type qos input pmap1 | Assign service-policy to interface on in-direction. |
| (config-if)#exit | Exit interface mode. |
| (config)#commit | Commit candidate configuration to the running configuration. |

P:NSM

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.10.2/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#ip address 10.10.11.1/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 192.168.0.2/32 secondary | Set the IP address of the loopback interface to 192.168.0.2/32. |
| (config-if)#commit | Commit the transaction. |

P:OSPF

| | |
|--------------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer |

| | |
|--|---|
| | identifying the routing process. |
| (config-router)#network 10.10.10.0/24 area 0 (config-router)#network 10.10.11.0/24 area 0 (config-router)#network 192.168.0.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

P:LDP

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 192.168.0.2 | Set the router ID to IP address 192.168.0.2. |
| (config-router)#transport-address ipv4 192.168.0.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| (config-if)#commit | Commit the transaction. |

PE2:NSM

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#ip address 10.10.11.2/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 192.168.0.3/32 secondary | Set the IP address of the loopback interface to 192.168.0.3/32. |
| (config-if)#commit | Commit the transaction. |

PE2:OSPF

| | |
|---------------------|-----------------------|
| #configure terminal | Enter configure mode. |
|---------------------|-----------------------|

| | |
|---|---|
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.10.11.0/24 area 0 (config-router)#network 192.168.0.3/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

PE2:LDP

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 192.168.0.3 | Set the router ID to IP address 192.168.0.3. |
| (config-router)#explicit-null | To disable PHP. |
| (config-router)#transport-address ipv4 192.168.0.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#enable-ldp ipv4 | Enable LDP on xe2. |
| (config-if)#commit | Commit the transaction. |

PE2

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#qos enable | Enable QoS. |
| (config)#hardware-profile filter qos enable | Enable QoS hardware. |
| (config)#commit | Commit candidate configuration to the running configuration. |
| (config)#class-map cmap1 | Enter Class-map mode. |
| (config-cmap-qos)#match dscp 2 | Configure match criteria as DSCP with Value 2. |
| (config-cmap-qos)#exit | Exit Class-map mode. |
| (config)#commit | Commit candidate configuration to the running configuration. |
| (config)#policy-map pmap1 | Enter policy-map mode. |
| (config-pmap-qos)#class cmap1 | Assign Class cmap1 to Policy-map pmap1. |
| (config-pmap-c-qos)#police cir 1 mbps eir 1 mbps bc 1 kbytes be 1 kbytes | Police DSCP 2 packets @ Committed information rate 1 mbps. |

| | |
|---|--|
| (config-pmap-c-qos)#exit | Exit out of policy-class-map mode. |
| (config-pmap-qos)#exit | Exit out of Policy-map mode. |
| (config)#commit | Commit candidate configuration to the running configuration. |
| (config)#interface xe1 | Enter xe1 interface. |
| (config-if)#service-policy type qos input pmap1 | Assign service-policy to interface on in-direction. |
| (config-if)#exit | Exit interface mode. |
| (config)#commit | Commit candidate configuration to the running configuration. |

Validation

PE1

```

PE1#show class-map

Type qos class-maps
=====
    class-map type qos match-any class-default

    class-map type qos match-all cmap1
        match dscp 2

Type queuing class-maps
=====
    class-map type queuing class-default-q

    class-map type queuing q0

    class-map type queuing q1

    class-map type queuing q2

    class-map type queuing q3

    class-map type queuing q4

    class-map type queuing q5

    class-map type queuing q6

    class-map type queuing q7

Type Vlan-Queuing class-maps
=====

#show running-config qos
qos enable
!
class-map type qos match-all cmap1

match dscp 2
!
policy-map type qos pmap1
class type qos cmap1
police cir 1 mbps eir 1 mbps bc 1 kbytes be 1 kbytes

```

```
exit
!
!
!
interface xel
  service-policy type qos input pmap1
#

#show policy-map

Type qos policy-maps
=====

policy-map type qospmap1
  class Type qos cmap1
    police cir 1 mbps eir 1 mbps bc 1 kbytes be 1 kbytes

Type queuing policy-maps
=====

policy-map type queuing default default-out-policy
  class type queuing default q0
    priority level 0
    exit
  class type queuing default q1
    priority level 1
    exit
  class type queuing default q2
    priority level 2
    exit
  class type queuing default q3
    priority level 3
    exit
  class type queuing default q4
    priority level 4
    exit
  class type queuing default q5
    priority level 5
    exit
  class type queuing default q6
    priority level 6
    exit
  class type queuing default q7
    priority level 7

policy-map type queuing default subif-default-out-policy
  class type queuing default q0
    priority level 0
  class type queuing default q1
    priority level 1
  class type queuing default q2
    priority level 2
  class type queuing default q3
    priority level 3

#show policy-map interface xel
=====

Interface xel

Type QoS statistics status : disabled
=====

Class-map (qos): cmap1 (match all)
  match dscp 2
  police cir 1 mbps eir 1 mbps bc 1 kbytes be 1 kbytes
```

```
Type Queuing policy-map : default-out-policy
=====

Class-map (queuing): q0
  shape 10000000 kbps (inherited)
  priority level 0
  queue-limit 12517376 bytes/10 ms (default)

Class-map (queuing): q1
  shape 10000000 kbps (inherited)
  priority level 1
  queue-limit 12517376 bytes/10 ms (default)

Class-map (queuing): q2
  shape 10000000 kbps (inherited)
  priority level 2
  queue-limit 12517376 bytes/10 ms (default)

Class-map (queuing): q3
  shape 10000000 kbps (inherited)
  priority level 3
  queue-limit 12517376 bytes/10 ms (default)

Class-map (queuing): q4
  shape 10000000 kbps (inherited)
  priority level 4
  queue-limit 12517376 bytes/10 ms (default)

Class-map (queuing): q5
  shape 10000000 kbps (inherited)
  priority level 5
  queue-limit 12517376 bytes/10 ms (default)

Class-map (queuing): q6
  shape 10000000 kbps (inherited)
  priority level 6
  queue-limit 12517376 bytes/10 ms (default)
  Output
    Total      : 7 packets, 816 bytes
    Green      : 7 packets, 816 bytes
    Yellow     : 0 packets, 0 bytes

Class-map (queuing): q7

shape 10000000 kbps (inherited)
priority level 7
queue-limit 12517376 bytes/10 ms (default)
```

PE2

```
OcNOS#show class-map

Type qos class-maps
=====
  class-map type qos match-any class-default

  class-map type qos match-all cmap1
    match dscp 2

Type queuing class-maps
=====
  class-map type queuing class-default-q

  class-map type queuing q0
```

```
class-map type queuing q1

class-map type queuing q2

class-map type queuing q3

class-map type queuing q4

class-map type queuing q5

class-map type queuing q6

class-map type queuing q7

OcNOS#show running-config qos
qos enable
!
class-map type qos match-all cmap1
match dscp 2
!
policy-map type qos pmap1
class type qos cmap1
  police cir 1 mbps eir 1 mbps bc 1 kbytes be 1 kbytes
exit
!
!
!
interface xel
  service-policy type qos input pmap1
!

OcNOS#show policy-map

Type qos policy-maps
=====

policy-map type qos pmap1
class type qos cmap1
  police cir 1 mbps eir 1 mbps bc 1 kbytes be 1 kbytes

Type queuing policy-maps
=====

policy-map type queuing default default-out-policy
class type queuing default q0
  priority level 0
class type queuing default q1
  priority level 1
class type queuing default q2
  priority level 2
class type queuing default q3
  priority level 3
class type queuing default q4
  priority level 4
class type queuing default q5
  priority level 5
class type queuing default q6
  priority level 6
class type queuing default q7
  priority level 7

policy-map type queuing default subif-default-out-policy
```

```
class type queuing default q0
  priority level 0
class type queuing default q1
  priority level 1
class type queuing default q2
  priority level 2
class type queuing default q3
  priority level 3

OcNOS#show policy-map int xe15

Interface xe15

Type QoS statistics status : disabled

Type QoS Ingress policy-map : pmap1

Class-map (qos): cmap1 (match all)
  match dscp 2
  police cir 1 mbps eir 1 mbps bc 1 kbytes be 1 kbytes

Type Queuing policy-map : default-out-policy

Class-map (queuing): q0
  shape 10000000 kbps (inherited)
  priority level 0
  queue-limit (default)
  Output
    Total      : 451 packets, 33887 bytes
    Green      : 451 packets, 33887 bytes
    Yellow     : 0 packets, 0 bytes

Class-map (queuing): q1
  shape 10000000 kbps (inherited)
  priority level 1
  queue-limit (default)

Class-map (queuing): q2
  shape 10000000 kbps (inherited)
  priority level 2
  queue-limit (default)

Class-map (queuing): q3
  shape 10000000 kbps (inherited)
  priority level 3
  queue-limit (default)

Class-map (queuing): q4
  shape 10000000 kbps (inherited)
  priority level 4
  queue-limit (default)

Class-map (queuing): q5
  shape 10000000 kbps (inherited)
  priority level 5
  queue-limit (default)

Class-map (queuing): q6
  shape 10000000 kbps (inherited)
  priority level 6
  queue-limit (default)
  Output
    Total      : 1119 packets, 100461 bytes
    Green      : 1119 packets, 100461 bytes
    Yellow     : 0 packets, 0 bytes

Class-map (queuing): q7
```



```
shape 10000000 kbps (inherited)
priority level 7
queue-limit (default)
  Output
    Total      : 56 packets, 4088 bytes
    Green      : 56 packets, 4088 bytes
    Yellow     : 0 packets, 0 bytes
```

MPLS Statistics Configuration

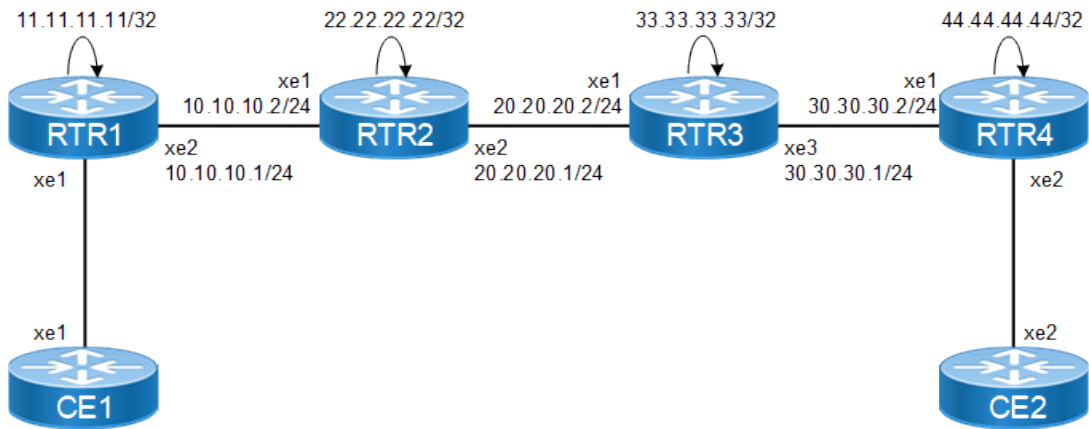
This chapter provides the configuration required for configuring MPLS LSPs and verifying the statistics of packets captured at the supported interfaces, in terms of both packet count and bytes, when traffic is sent.



Note: Starting with OcNOS version 7.0.0, sub-interfaces replace Service Templates as the preferred configuration method. While legacy Service Template configurations remain functional, the associated commands are now hidden and will be completely removed in a future release. To ensure future-proofing and continued manageability, please migrate all existing Service Templates to sub-interfaces. For more details, see [Appendix A: Service Template Configuration \(page 2631\)](#).

Topology

Figure 42. MPLS Statistics Topology



Configure LDP-LSP

RTR1: Loopback Interface configuration

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the interface (lo) to be configured. |
| (config-if)#ip address 11.11.11.11/32 secondary | Configure IP address on loopback interface |
| (config-if)#commit | Commit the transaction. |

RTR1: Global LDP configuration

| | |
|--|--|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 11.11.11.11 | Configure the router-id |
| (config-router)#transport-address ipv4 11.11.11.11 | Configure the loopback address as transport- |

| | |
|--|---|
| | address |
| (config-router)#targeted-peer ipv4 44.44.44.44 | Configure the loopback address of RTR4 as targeted peer. |
| (config-router-targeted-peer)#exit | Exit router-targeted-peer mode and enter config-router mode |
| (config-router-targeted-peer)#commit | Commit the transaction. |

RTR1: Enabling LDP and label switching on interface

| | |
|--------------------------------------|--|
| (config)#interface xe2 | Enter interface mode for xe2. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#ip address 10.10.10.1/24 | Configure IP address on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR1: Global OSPF configuration

| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#network 11.11.11.11/32 area 0 | Advertise loopback address in OSPF. |
| (config-router)#network 10.10.10.0/24 area 0 | Advertise network address (xe2) in OSPF. |
| (config-router)#commit | Commit the transaction. |

RTR2: Loopback Interface configuration

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the interface (lo) to be configured. |
| (config-if)#ip address 22.22.22.22/32 secondary | Configure IP address on loopback interface |
| (config-if)#commit | Commit the transaction. |

RTR2: Global LDP configuration

| | |
|--|---|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 22.22.22.22 | Configure the router-id |
| (config-router)#transport-address ipv4 22.22.22.22 | Configure the loopback address as transport-address |
| (config-router)#commit | Commit the transaction. |

RTR2: Enabling LDP and label switching on interface

| | |
|-----------------------------|-------------------------------|
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the interface. |

| | |
|--------------------------------------|--|
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#ip address 10.10.10.2/24 | Configure IP address on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode for xe2. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#ip address 20.20.20.1/24 | Configure IP address on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR2: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode |
| (config-router)#network 22.22.22.22/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.10.10.2/24 area 0.0.0.0 | Advertise network address (xe1) in OSPF. |
| (config-router)#network 20.20.20.1/24 area 0.0.0.0 | Advertise network address (xe2) in OSPF. |
| (config-router)#commit | Commit the transaction. |

RTR3: Loopback Interface configuration

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the interface (lo) to be configured. |
| (config-if)#ip address 33.33.33.33/32 secondary | Configure IP address on loopback interface |
| (config-if)#commit | Commit the transaction. |

RTR3: Global LDP configuration

| | |
|--|---|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 33.33.33.33 | Configure the router-id |
| (config-router)#transport-address ipv4 33.33.33.33 | Configure the loopback address as transport-address |
| (config-router)#commit | Commit the transaction. |

RTR3: Enabling LDP and label switching on interface

| | |
|--------------------------------------|--|
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#ip address 20.20.20.2/24 | Configure IP address on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode for xe2. |

| | |
|--------------------------------------|--|
| (config-if)#enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#ip address 30.30.30.1/24 | Configure IP address on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR3: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode |
| (config-router)#network 33.33.33.33/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 20.20.20.2/24 area 0.0.0.0 | Advertise network address (xe1) in OSPF. |
| (config-router)#network 30.30.30.1/24 area 0.0.0.0 | Advertise network address (xe2) in OSPF. |
| (config-router)#commit | Commit the transaction. |

RTR4: Loopback Interface configuration

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the interface (lo) to be configured. |
| (config-if)#ip address 44.44.44.44/32 secondary | Configure IP address on loopback interface |
| (config-if)#commit | Commit the transaction. |

RTR4: Global LDP configuration

| | |
|--|---|
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 44.44.44.44 | Configure the router-id |
| (config-router)#transport-address ipv4 44.44.44.44 | Configure the loopback address as transport-address |
| (config-router)#targeted-peer ipv4 11.11.11.11 | Configure the loopback address of RTR1 as targeted peer. |
| (config-router-targeted-peer)#exit | Exit router-targeted-peer mode and enter config-router mode |
| (config-router)#commit | Commit the transaction. |

RTR4: Enabling LDP and label switching on interface

| | |
|--------------------------------------|--|
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#ip address 30.30.30.2/24 | Configure IP address on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR4: Global OSPF Configuration

| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode |
| (config-router)#network 44.44.44.44/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 30.30.30.2/24 area 0.0.0.0 | Advertise network address (xe1) in OSPF. |
| (config-router)#commit | Commit the transaction. |

MPLS LDP VPLS Configuration**RTR1: VPLS**

| | |
|---|--|
| (config)#mpls vpls vpls1 1 | Enter the VPLS configuration mode |
| (config-vpls)# signaling ldp | Use LDP signaling for VPLS |
| (config-vpls-sig)#vpls-type ethernet | Configure the VPLS as Ethernet |
| (config-vpls-sig)#vpls-peer 44.44.44.44 | Configure RTR4 as VPLS peer for RTR1. |
| (config-vpls-sig)#exit-signaling | Exit signaling mode |
| (config-vpls)# exit-vpls | Exit VPLS config mode and return to Configure mode |
| (config-vpls)#commit | Commit the transaction. |

RTR1: Interface

| | |
|--|--|
| (config)#service-template st1 | Template configuration |
| (config-svc)#exit | Exit service template mode |
| (config)#interface xe1 | Enter the Interface mode for xe1 |
| (config-if)#switchport | Enable switchport on the interface |
| (config-if)#mpls-vpls vpls1 service-template st1 | Bind the VPLS to the Access Interface. |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if)#commit | Commit the transaction. |

RTR4: VPLS

| | |
|---|---------------------------------------|
| (config)#mpls vpls vpls1 1 | Enter the VPLS configuration mode |
| (config-vpls)#signaling ldp | Use LDP signaling for VPLS |
| (config-vpls-sig)#vpls-type ethernet | Configure the VPLS as Ethernet |
| (config-vpls-sig)#vpls-peer 11.11.11.11 | Configure RTR1 as VPLS peer for RTR4. |
| (config-vpls-sig)#exit-signaling | Exit signaling mode |
| (config-vpls)#commit | Commit the transaction. |

RTR4: Interface

| | |
|-------------------------------|------------------------|
| (config)#service-template st1 | Template configuration |
|-------------------------------|------------------------|

| | |
|---|--|
| <code>(config-svc)#exit</code> | Exit service template mode |
| <code>(config)#interface xe2</code> | Enter the Interface mode for xe2 |
| <code>(config-if)#switchport</code> | Enable switchport on the interface |
| <code>(config-if)#mpls-vpls vpls1 service template st1</code> | Bind the VPLS to the Access Interface. |
| <code>(config-if-vpls)#split-horizon group access1</code> | Configure split-horizon group on VPLS |
| <code>(config-if-vpls)#commit</code> | Commit the transaction. |

Virtual Circuit Configuration

RTR1: Global VC Configuration

| | |
|--|--|
| <code>(config)#mpls l2-circuit t1 100 44.44.44.44</code> | Enter the VC configuration command in router mode. |
| <code>(config-pseudowire)#exit</code> | Exit service template mode |
| <code>(config)#bridge 1 protocol ieee vlan-bridge</code> | Creating a VLAN-bridge in router mode. |
| <code>(config)#commit</code> | Commit the transaction. |

RTR1: Interface Configuration

| | |
|---|---------------------------------------|
| <code>(config)#service-template st1</code> | Template configuration |
| <code>(config-svc)#exit</code> | Exit service template configuration |
| <code>(config)#interface xe1</code> | Enter interface mode for xe1. |
| <code>(config-if)#switchport</code> | Enable switchport on the interface. |
| <code>(config-if)#mpls-vpls vpls1 service-template st1</code> | Bind the interface to VPLS. |
| <code>(config-if-vpls)#split-horizon group access1</code> | Configure split-horizon group on VPLS |
| <code>(config-if)#commit</code> | Commit the transaction. |

RTR4: Global VC Configuration

| | |
|--|--|
| <code>(config)#mpls l2-circuit t1 100 11.11.11.11</code> | Enter the VC configuration command in router mode. |
| <code>(config)#exit</code> | Exit service template mode. |
| <code>(config)#bridge 1 protocol ieee vlan-bridge</code> | Creating a VLAN-bridge in router mode. |
| <code>(config)#commit</code> | Commit the transaction. |

RTR4: Interface Configuration

| | |
|--|-------------------------------------|
| <code>(config)#service-template st1</code> | Template configuration |
| <code>(config-svc)#exit</code> | Exit service template mode |
| <code>(config)#interface xe2</code> | Enter interface mode for xe2. |
| <code>(config-if)#switchport</code> | Enable switchport on the interface. |

| | |
|--|---------------------------------------|
| (config-if)#mpls-vpls vpls1 service-template st1 | Bind the interface to VPLS. |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config)#commit | Commit the transaction. |

VPLS Configuration

RTR1: Global VPLS Configuration

| | |
|---|---------------------------------------|
| (config)#mpls vpls vpls1 1 | Enter the VPLS configuration mode. |
| (config-vpls)#signaling ldp | Use LDP signaling for VPLS. |
| (config-vpls-sig)#vpls-peer 44.44.44.44 | Configure RTR4 as VPLS peer for RTR1. |
| (config-vpls-sig)#exit-signaling | Exit signaling mode. |
| (config-vpls)#commit | Commit the transaction. |

RTR1: Interface Configuration

| | |
|--|-------------------------------------|
| (config)#service-template st1 | Template configuration |
| (config-svc)#exit | Exit service template mode |
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#switchport | Enable switchport on the interface. |
| (config-if)#mpls-vpls vpls1 service-template st1 | Bind the interface to VPLS. |
| (config-if-vpls)#commit | Commit the transaction. |

RTR4: Global VC Configuration

| | |
|---|---------------------------------------|
| (config)#mpls vpls vpls1 1 | Enter the VPLS configuration mode. |
| (config-vpls)#signaling ldp | Use LDP signaling for VPLS. |
| (config-vpls-sig)#vpls-peer 11.11.11.11 | Configure RTR4 as VPLS peer for RTR1. |
| (config-vpls-sig)#exit-signaling | Exit signaling mode. |
| (config-vpls)#commit | Commit the transaction. |

RTR4: Interface Configuration

| | |
|--|--|
| (config)#service-template st1 | Template configuration |
| (config-svc)#exit | Exit service template mode |
| (config)#interface xe2 | Enter interface mode for xe2. |
| (config-if)#switchport | Enable switchport on the interface. |
| (config-if)#mpls-vpls vpls1 service-template st1 | Bind the interface to VPLS. |
| (config-if-vpls)#exit | Exit VPLS mode and return to interface mode. |
| (config-if)#commit | Commit the transaction. |

Configure RSVP-LSP

RTR1: Global RSVP configuration

| | |
|------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#router rsvp | Enter RSVP configuration mode for the router. |
| (config-router)#commit | Commit the transaction. |

RTR1: Enabling RSVP and label switching on interface

| | |
|-----------------------------|--|
| (config)#interface xe2 | Enter interface mode for xe2. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR1: Trunk Configuration

| | |
|-------------------------------|------------------------------------|
| (config)#rsvp-trunk t1 | Configure RSVP trunk t1 |
| (config-trunk)#to 44.44.44.44 | Configure RTR4 as the end of trunk |
| (config-trunk)#commit | Commit the transaction. |

RTR2: Global RSVP configuration

| | |
|------------------------|---|
| (config)#router rsvp | Enter RSVP configuration mode for the router. |
| (config-router)#php | Configure PHP on the end node. |
| (config-router)#commit | Commit the transaction. |

RTR2: Enabling RSVP and label switching on interface

| | |
|-----------------------------|--|
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode for xe2. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR3: Global RSVP configuration

| | |
|---------------------|-----------------------|
| #configure terminal | Enter configure mode. |
|---------------------|-----------------------|

| | |
|------------------------|---|
| (config)#router rsvp | Enter RSVP configuration mode for the router. |
| (config-router)#php | Configure PHP on the end node. |
| (config-router)#commit | Commit the transaction. |

RTR3: Enabling RSVP and label switching on interface

| | |
|-----------------------------|--|
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe3 | Enter interface mode for xe3. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR4: Global RSVP configuration

| | |
|------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#router rsvp | Enter RSVP configuration mode for the router. |
| (config-router)#commit | Commit the transaction. |

RTR4: Enabling RSVP and label switching on interface

| | |
|-----------------------------|--|
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR4: Trunk Configuration

| | |
|-------------------------------|-------------------------------------|
| (config)#rsvp-trunk t2 | Configure RSVP trunk t2. |
| (config-trunk)#to 11.11.11.11 | Configure RTR1 as the end of trunk. |
| (config-trunk)#commit | Commit the transaction. |

Configure Static-LSP

RTR1: Global Static configuration

| | |
|---|--|
| (config)#mpls ftn-entry 44.44.44.44/32 100 10.10.10.2 xe1 | Configure FTN entry for rtr4 loopback. |
| (config)#mpls ilm-entry 900 pop | Pop the incoming label |
| (config)#commit | Commit the transaction. |

RTR1: Enabling label switching on interface

| | |
|-----------------------------|--|
| (config)#interface xe2 | Enter interface mode for xe2. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config)#commit | Commit the transaction. |

RTR2: Global Static configuration

| | |
|--|-------------------------|
| mpls ilm-entry 100 swap 200 xe2 20.20.20.2 44.44.44.44/32 | Swap the incoming label |
| mpls ilm-entry 800 swap 900 xe1 10.10.10.1 11.11.11.11/32 | Swap the incoming label |

RTR2: Enabling label switching on interface

| | |
|-----------------------------|--|
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode for xe2. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR3: Global Static configuration

| | |
|---|--------------------------|
| (config)#mpls ilm-entry 200 swap 300 xe2 30.30.30.2 44.44.44.44/32 | Swap the incoming label. |
| (config)#mpls ilm-entry 700 swap 800 xe1 20.20.20.1 11.11.11.11/32 | Swap the incoming label. |
| (config)#commit | Commit the transaction. |

RTR3: Enabling label switching on interface

| | |
|-----------------------------|--|
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe3 | Enter interface mode for xe3. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#commit | Commit the transaction. |

RTR4: Global Static configuration

| | |
|---|--|
| (config)#mpls ftn-entry 11.11.11.11/32 700 30.30.30.1 | Configure FTN entry for RTR1 loopback. |
|---|--|

| | |
|--------------------------------|-------------------------|
| xe1 | |
| (config)mpls ilm-entry 300 pop | Pop the incoming label. |
| (config)#commit | Commit the transaction. |

RTR4: Enabling label switching on interface

| | |
|-----------------------------|--|
| (config)#interface xe1 | Enter interface mode for xe1. |
| (config-if)#label-switching | Enable Label switching on the interface. |
| (config-if)#commit | Commit the transaction. |

Validation

Here, 1000 packets are transmitted between the PE nodes and the output of counters at each node is mentioned below.

For Static-LSP

```
RTR1#show mpls counters static
```

```
[FTN statistics]
```

```
+-----+-----+-----+-----+
|      FEC      | out-label | Tx packets | Tx bytes |
+-----+-----+-----+-----+
| 44.44.44.44/32 | 100       | 49939      | 807798   |
```

```
[ILM statistics]
```

```
+-----+-----+-----+-----+-----+-----+
|      FEC      | in-label | out-label | Rx packets | Rx bytes | Tx
| packets      | Tx bytes |           |            |          |
+-----+-----+-----+-----+-----+-----+
| 0.0.0.0/0     | 900      | n/a       | 40546      | 3486956  | n/a
| n/a           |          |           |            |          |
```

```
RTR1#
```

```
RTR2#show mpls counters static
```

```
[FTN statistics]
```

```
+-----+-----+-----+-----+
|      FEC      | out-label | Tx packets | Tx bytes |
+-----+-----+-----+-----+
```

```
[ILM statistics]
```

```
+-----+-----+-----+-----+-----+-----+
|      FEC      | in-label | out-label | Rx packets | Rx bytes |
| Tx packets    | Tx bytes |           |            |          |
+-----+-----+-----+-----+-----+-----+
| 44.44.44.44/32 | 100      | 200       | 9393       | 807798   |
| 9393           |          | 807798    |            |          |
| 11.11.11.11/32 | 800      | 900       | 40546      | 3486956  |
| 40546          |          | 3486956   |            |          |
```

```
RTR2#
```

```
RTR3#show mpls counters static
```

```
[FTN statistics]
```

```
+-----+-----+-----+-----+
|      FEC      | out-label | Tx packets | Tx bytes |
+-----+-----+-----+-----+
```

```
[ILM statistics]
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
|      FEC      | in-label | out-label | Rx packets | Rx bytes | Tx
| packets      | Tx bytes |          |            |          |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 44.44.44.44/32 | 200      | 300      | 9393       | 807798   | 9393
| 807798         |          |          |            |          |
| 11.11.11.11/32 | 700      | 800      | 40546      | 3486956  | 40546
| 3486956        |          |          |            |          |
RTR3#

RTR4#show mpls counters static
[FTN statistics]
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
|      FEC      | out-label | Tx packets | Tx bytes |
+-----+-----+-----+-----+-----+
| 11.11.11.11/32 | 700       | 49939     | 3486956  |
+-----+-----+-----+-----+-----+

[ILM statistics]
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
|      FEC      | in-label | out-label | Rx packets | Rx bytes |
| Tx packets    | Tx bytes |          |            |          |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 0.0.0.0/0     | 300      | n/a       | 9393      | 807798   |
| n/a           |          | n/a       |            |          |
RTR4#
```

For RSVP-LSP

```
RTR1#show mpls counters rsvp
Tunnel-id 5001 Extended Tunnel-ID 44.44.44.44 Egress 11.11.11.11
lsp-name : t1-Primary [Egress]
lsp-ingress : 44.44.44.44 lsp-id : 2201
Rx pkts : 2509072 Rx bytes : 187780192
Tx pkts : 0 Tx bytes : 0

Tunnel-id 5001 Extended Tunnel-ID 11.11.11.11 Egress 44.44.44.44
lsp-name : t1-Primary [Ingress]
lsp-ingress : 11.11.11.11 lsp-id : 2201
Rx pkts : 0 Rx bytes : 0
Tx pkts : 5578405 Tx bytes : 451417492

RTR2#show mpls counters rsvp
Tunnel-id 5001 Extended Tunnel-ID 44.44.44.44 Egress 11.11.11.11
lsp-name : t1-Primary [Transit]
lsp-ingress : 44.44.44.44 lsp-id : 2201
Rx pkts : 2565947 Rx bytes : 192671442
Tx pkts : 2565960 Tx bytes : 192672560

Tunnel-id 5001 Extended Tunnel-ID 11.11.11.11 Egress 44.44.44.44
lsp-name : t1-Primary [Transit]
lsp-ingress : 11.11.11.11 lsp-id : 2201
Rx pkts : 5631460 Rx bytes : 456305560
Tx pkts : 5631472 Tx bytes : 456306592

RTR3#show mpls counters rsvp
Tunnel-id 5001 Extended Tunnel-ID 44.44.44.44 Egress 11.11.11.11
lsp-name : t1-Primary [Transit]
lsp-ingress : 44.44.44.44 lsp-id : 2201
```

```

Rx pkts : 2565947          Rx bytes : 282671442
Tx pkts : 2565960          Tx bytes : 282672560

Tunnel-id 5001 Extended Tunnel-ID 11.11.11.11 Egress 44.44.44.44
  lsp-name : t1-Primary          [Transit]
  lsp-ingress : 11.11.11.11      lsp-id : 2201
Rx pkts : 8631460          Rx bytes : 457245560
Tx pkts : 8631472          Tx bytes : 45724592

RTR4#show mpls counters rsvp
Tunnel-id 5001 Extended Tunnel-ID 44.44.44.44 Egress 11.11.11.11
  lsp-name : t1-Primary          [Ingress]
  lsp-ingress : 44.44.44.44      lsp-id : 2201
Rx pkts : 0                Rx bytes : 0
Tx pkts : 10231330         Tx bytes : 374371318

Tunnel-id 5001 Extended Tunnel-ID 11.11.11.11 Egress 44.44.44.44
  lsp-name : t1-Primary          [Egress]
  lsp-ingress : 11.11.11.11      lsp-id : 2201
Rx pkts : 5651207          Rx bytes : 458003802
Tx pkts : 0                Tx bytes : 0

R3#show mpls counters rsvp
Tunnel-id 5001 Extended Tunnel-ID 44.44.44.44 Egress 11.11.11.11 lsp-name : t1-Primary  [Transit]
lsp-ingress : 44.44.44.44  lsp-id : 2201
Rx pkts : 2565947  Rx bytes : 282671442
Tx pkts : 2565960  Tx bytes : 282672560

Tunnel-id 5001 Extended Tunnel-ID 11.11.11.11 Egress 44.44.44.44 lsp-name : t1-Primary  [Transit]
lsp-ingress : 11.11.11.11  lsp-id : 2201
Rx pkts : 8631460  Rx bytes : 457245560
Tx pkts : 8631472  Tx bytes : 45724592

R4#show mpls counters rsvp
Tunnel-id 5001 Extended Tunnel-ID 44.44.44.44 Egress 11.11.11.11 lsp-name : t1-Primary  [Ingress]
lsp-ingress : 44.44.44.44  lsp-id : 2201 Rx pkts : 0  Rx bytes : 0
Tx pkts : 10231330  Tx bytes : 374371318

Tunnel-id 5001 Extended Tunnel-ID 11.11.11.11 Egress 44.44.44.44 lsp-name : t1-Primary  [Egress]
lsp-ingress : 11.11.11.11  lsp-id : 2201
Rx pkts : 5651207  Rx bytes : 458003802
Tx pkts : 0  Tx bytes : 0

```

For LDP-LSP

```

RTR1#show mpls counters ldp
[FTN statistics]
+-----+-----+-----+-----+
|      FEC      | out-label | Tx packets | Tx bytes |
+-----+-----+-----+-----+
| 44.44.44.44/32 | 52483     | 1000       | 1004000  |
+-----+-----+-----+-----+
[ILM statistics]
+-----+-----+-----+-----+-----+-----+-----+
| FEC | in-label | out-label | Rx packets | Rx bytes | Tx packets | Tx bytes |
+-----+-----+-----+-----+-----+-----+-----+

RTR2#show mpls counters ldp
[FTN statistics]
+-----+-----+-----+-----+
|      FEC      | out-label | Tx packets | Tx bytes |
+-----+-----+-----+-----+
[ILM statistics]
+-----+-----+-----+-----+-----+-----+-----+
| FEC | in-label | out-label | Rx packets | Rx bytes | Tx packets | Tx bytes |
+-----+-----+-----+-----+-----+-----+-----+

```

```
44.44.44.44/32      52483    52483  1000      1004000    1000      1004000
```

```
/RTR2#show mpls counters ldp [FTN statistics]
```

```
+-----+-----+-----+-----+
|      FEC      | out-label |   Tx packets   |   Tx bytes   |
+-----+-----+-----+-----+
[ILM statistics]
+-----+-----+-----+-----+-----+-----+-----+
| FEC | in-label | out-label | Rx packets | Rx bytes | Tx packets | Tx bytes |
+-----+-----+-----+-----+-----+-----+-----+
44.44.44.44/32  52483    52483  1000    1004000    1000    1004000
```

For LDP-VC

```
RTR1#show mpls l2-circuit t1 statistics
MPLS Layer-2 Virtual Circuit: t1, id 100
```

```
Access port statistics:
```

```
  RX:  Input packets : 0
       Input bytes   : 0
  TX:  Output packets : 4642811
       Output bytes  : 297139904
```

```
Network port statistics:
```

```
  RX:  Input packets : 4642804
       Input bytes   : 399281144
  TX:  Output packets : 0
       Output bytes  : 0
```

```
RTR4#show mpls l2-circuit t1 statistics
MPLS Layer-2 Virtual Circuit: t1, id 100
```

```
Access port statistics:
```

```
  RX:  Input packets : 4633957
       Input bytes   : 296573248
  TX:  Output packets : 0
       Output bytes  : 0
```

```
Network port statistics:
```

```
  RX:  Input packets : 0
       Input bytes   : 0
  TX:  Output packets : 4633960
       Output bytes  : 398520560
```

For LDP-VPLS

```
RTR1#show mpls vpls vpls1 statistics
Virtual Private LAN Service Instance: vpls1, ID: 1
```

```
Access port statistics:
```

```
Interface: xel
```

```
  RX:  Input packets : 1922483
       Input bytes   : 123038912
  TX:  Output packets : 3894242
       Output bytes  : 126192000
```

```
Network port statistics:
```

```
Mesh Peer: 44.44.44.44 (Up)
```

```
  RX:  Input packets : 1971746
       Input bytes   : 169570156
  TX:  Output packets : 3894244
       Output bytes  : 165334398
```

```
RTR4#show mpls vpls vpls1 statistics
```

Virtual Private LAN Service Instance: vpls1, ID: 1

Access port statistics:

Interface: xe2

| | | |
|-----|----------------|-------------|
| RX: | Input packets | : 1967571 |
| | Input bytes | : 125924544 |
| TX: | Output packets | : 3885889 |
| | Output bytes | : 122772032 |

Network port statistics:

Mesh Peer: 11.11.11.11 (Up)

| | | |
|-----|----------------|-------------|
| RX: | Input packets | : 1918310 |
| | Input bytes | : 164974660 |
| TX: | Output packets | : 3885892 |
| | Output bytes | : 169211622 |

MPLS Service Mapping Configuration

This chapter includes step-by-step configurations for tunnel policy configuration for all the MPLS services.

Overview

This feature enables MPLS services L2VPN, L3VPN, EVPN traffic steering with SR-policy as the underlying transport using local-tunnel-policy configuration. This is achieved with the help of color coding and association between service and MPLS transport.

A local-tunnel-policy is created which specifies the color of the transport-tunnel in order to be selected for a VPN service. The created tunnel-policy is now applied for the VPN service so that it will select the preferred tunnel. The destination address of the preferred tunnel identifies a peer PE so the PE will forward traffic destined for that peer PE over the preferred tunnel. If you specify multiple preferred tunnels that have the same destination address in a tunnel policy, only the first configured tunnel takes effect and if the first tunnel is not available, the second tunnel is used, and so forth. No load balancing will be performed on these tunnels.


The tunnels selected by 2 methods, dedicated and best-effort. By default, tunnel policy uses dedicated method. If you configure dedicated method for a tunnel policy, the tunnel policy uses the preferred tunnel to forward traffic destined for the peer PE. If not, the MPLS VPN service will be down. If you configure best-effort method for a tunnel policy, the tunnel policy uses the preferred tunnel to forward traffic destined for the peer PE. If not, the MPLS VPN selects the available MPLS transport in MPLS forwarding table.



Note: When a tunnel policy configuration is attached to a VRF, the policy is applied to all the routes learned within that VRF. In an ECMP scenario, it is expected that all the ECMP paths are associated with an SR policy as defined by the configured `tunnel-policy`. Due to this, the use of service mapping is not ideal in ECMP. Instead, an SR-ODN or color-based mapping is recommended as a more suitable alternative.

Topology

The diagram depicts the topology for the configuration examples that follow.

Figure 43. Service mapping to Tunnel

Configuration



Note: Starting with OcNOS version 7.0.0, sub-interfaces replace Service Templates as the preferred configuration method. While legacy Service Template configurations remain functional, the associated commands are now hidden and will be completely removed in a future release. To ensure future-proofing and continued manageability, please migrate all existing Service Templates to sub-interfaces. For more details, see [Appendix A: Service Template Configuration \(page 2631\)](#).

Below are the sample configuration for EVPN (E-LINE, E-LAN), L3VPN, 6VPE, VPWS, and VPLS services with Tunnel Policy along with SR as transport.

PE1: Loopback Interface

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.1/32 secondary | Configure IP address on loopback interface. |
| (config-if)#prefix-sid index 1 | Configure prefix sid index value |
| (config-if)#exit | Exit interface mode |

PE1: Global EVPN MPLS Command:

| | |
|--|---|
| #configure terminal | Enter configuration mode. |
| (config)#evpn mpls enable | Enable EVPN MPLS. Note: Reload is required after Enabling/ Disabling EVPN MPLS Feature. |
| (config)#evpn mpls vtep-ip-global 10.143.73.1 | Configuring loopback IP as VTEP global IP |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | Ending config mode |
| #write memory | Save config to startup config |
| #reload | Reload device after enabling evpn |

PE1: LDP Configuration

| | |
|---|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.1 | Set the router ID to IP address 10.143.73.1. |
| (config-router)#transport-address ipv4 10.143.73.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |

| | |
|--|---|
| (config-router)# targeted-peer ipv4 10.143.73.3 | Configure targeted peer |
| (config-router-targeted-peer)# exit-targeted-peer-mode | Exit from router target peer and LDP mode |

PE1: Interface Configuration Network Side

| | |
|--|--|
| (config)#interface xe16 | Enter the Interface mode for xe16 |
| (config-if)# ip address 10.255.128.2/31 | Configure IP address on the interface. |
| (config-if)#mtu 9900 | Configure mtu at interface level |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)# enable-ldp ipv4 | Enable ldp on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe4 | Enter the Interface mode for xe4. |
| (config-if)# ip address 10.255.128.10/31 | Configure IP address on the interface. |
| (config-if)#mtu 9900 | Configure mtu at interface level |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)# enable-ldp ipv4 | Enable ldp on the interface. |
| (config-if)#exit | Exit interface mode |

PE1: OSPF Configuration

| | |
|--|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.1 | Router-id configurations |
| (config-router)# bfd all-interfaces | Enable the OSPF enabled interfaces with bfd |
| (config-router)# network 10.143.73.1/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)# network 10.255.128.2/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# network 10.255.128.10/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# ospf segment-routing global block 16000 17000 | Enable SRGB under ospf process 100 |
| (config-router)#segment-routing mpls | Enable segment routing |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |

PE1: BGP Configuration

| | |
|--|--|
| (config)# router bgp 100 | Enter the Router BGP mode, ASN: 100. |
| (config-router)# neighbor 10.143.73.3 remote-as 100 | Configuring PE3 as iBGP neighbor using it's loopback ip. |
| (config-router)# neighbor 10.143.73.3 fall-over bfd multihop | Configure neighbor bfd multihop |
| (config-router)# neighbor 10.143.73.3 update-source lo | Source of routing updates as loopback |
| (config-router)# neighbor 10.143.73.3 advertisement-interval 0 | Configure advertisement interval |
| (config-router)# address-family vpnv4 unicast | Enter VPNv4 Address family mode |
| (config-router-af)# neighbor 10.143.73.3 activate | Enabling VPNv4 Address family for neighbor.. |
| (config-router-af)# exit-address-family | Exit Address-family mode. |
| (config-router)# address-family l2vpn vpls | Enter l2vpn Address family mode |
| (config-router-af)# neighbor 10.143.73.3 activate | Enabling l2vpn Address family for neighbor.. |
| (config-router-af)# exit-address-family | Exit Address-family mode. |
| (config-router)# address-family l2vpn evpn | Enter evpn Address family mode |
| (config-router-af)# neighbor 10.143.73.3 activate | Enabling evpn Address family for neighbor.. |
| (config-router-af)# exit-address-family | Exit Address-family mode. |
| (config-router)# address-family vpnv6 unicast | Enter VPNv6 Address family mode |
| (config-router-af)# neighbor 10.143.73.3 activate | Enabling VPNv6 Address family for neighbor.. |
| (config-router-af)# exit-address-family | Exit Address-family mode. |

PE1: SR policy Configuration

| | |
|--|---|
| (config)#segment-routing | Enter the segment routing mode. |
| (config-sr)# mpls sr-prefer | prefer segment routing MPLS entries for forwarding. |
| ((config-sr)#traffic-engineering | Enter traffic engineering mode. |
| (config-sr-te)# policy PE1-P1-PE3 | Configure SR policy |
| (config-sr-pol)# color 1 end-point 10.143.73.3 | Configure color with destination address. |

| | |
|---|---|
| (config-sr-pol)#candidate-path 1 | Enter candidate path mode. |
| (config-sr-pol-cp)#dynamic-path ospf 100 | Configure dynamic path ospf |
| (config-sr-pol-cp)#constraints | Enter constraint mode |
| (config-sr-dyn-cp-cons)# 10.143.73.5 loose | Configure P1 as next-hop loose mode. |
| (config-sr-dyn-cp-cons)#exit-pol-cp | Exit candidate path mode. |
| (config-sr-pol-cp)#exit-sr-pol | Exit SR policy mode. |
| (config-sr-te)# policy PE1-P2-PE3 | Configure SR policy |
| (config-sr-pol)# color 2 end-point 10.143.73.3 | Configure color with destination address. |
| (config-sr-pol)#candidate-path 1 | Enter candidate path mode. |
| (config-sr-pol-cp)#dynamic-path ospf 100 | Configure dynamic path ospf |
| (config-sr-pol-cp)#constraints | Enter constraint mode |
| (config-sr-dyn-cp-cons)# 10.143.73.6 loose | Configure P2 as next-hop loose mode. |
| (config-sr-dyn-cp-cons)#exit-pol-cp | Exit candidate path mode. |
| (config-sr-pol-cp)#exit-sr-pol | Exit SR policy mode. |

PE1: Tunnel Policy Configuration

| | |
|------------------------------------|---|
| (config)# tunnel-policy PE1-P1-PE3 | Configure tunnel policy |
| (config-tnl-policy)#color 1 | Map the color with the SR policy to be used.. |
| (config-tnl-policy)#exit | Exiting tunnel policy mode |
| (config)# tunnel-policy PE1-P2-PE3 | Configure tunnel policy |
| (config-tnl-policy)#color 2 | Map the color with the SR policy to be used.. |
| (config-tnl-policy)#exit | Exiting tunnel policy mode |

PE1: VPWS Configuration

| | |
|--|---|
| (config)# mpls l2-circuit pe1-to-pe3 2002 10.143.73.3 | Configure the VC for PE3. In this example, pe1-to-pe3 is the VC name, 2002 is the VC ID, and 10.143.73.3 is the VC endpoint IP address. |
| (config-pseudowire)# tunnel-select-policy PE1-P1-PE3 | Configure Tunnel-Policy on VC pe1-to-pe3 |
| (config-pseudowire)#exit | Exit pseudowire config mode. |
| (config)#exit | Exiting from config mode |
| (config)#service-template pe1-to-pe3 | Template configuration. |

| | |
|-------------------------------------|---|
| (config-svc)# match outer-vlan 2002 | Match criteria under template configuration |
| (config)#exit | Exiting from config mode |

PE1: VPLS Configuration

| | |
|---|--|
| (config)# mpls vpls vpls-pe1-to-pe3pe4-1 2502 | Enter VPLS config mode |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type ethernet | Type ethernet configuration for VPLS |
| (config-vpls-sig)# vpls-peer 10.143.73.3 tunnel-select-policy PE1-P2-PE3 | Configure VPLS Peer with tunnel-policy |
| (config-vpls-sig)#exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#service-template vpls-pe1-to-pe3pe4-1 | Template configuration. |
| (config-svc)# match outer-vlan 2502 | Match criteria under template configuration |
| (config-svc)# rewrite ingress translate 2600 outgoing-tpid dot1.q | Action performed for service template. |
| (config)# mpls vpls vpls-pe1-to-pe3pe4-2 2503 | Enter VPLS config mode |
| (config-vpls)#signaling bgp | Enter the Signaling bgp mode for BGP VPLS. |
| (config-vpls-sig)# ve-id 3 | Configure VE ID, which is mandatory for BGP VPLS, otherwise, signaling does not take place. VE ID should be unique per VPLS instance |
| (config-vpls-sig)# tunnel-select-policy PE1-P1-PE3 | Configure tunnel-policy for a VPLS instance |
| (config-vpls-sig)#exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#service-template vpls-pe1-to-pe3pe4-2 | Template configuration. |
| (config-svc)# match double-tag outer-vlan 2503 inner-vlan 2504 | Match criteria under template configuration |
| (config-svc)# rewrite ingress push 2505 | Action performed for service template. |

PE1: MAC VRF and IP VRF Configuration for EVPN and L3VPN services



Note: For EVPN Service, Tunnel Policy can be applied at mac vrf level and VNID level, When Configured at both levels, Preference is given to VNID level.

| | |
|--|---|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)# tunnel-select-policy PE1-PE3 | Map tunnel policy to mac vrf |
| (config-vrf)#rd 10.143.73.1:2 | Configuring Route-Distinguisher value 10.143.73.1:2 |
| (config-vrf)#route-target both 2:2 | Configuring import and export value as 2:2 |
| (config-vrf)#exit | Exiting VRF Mode. |
| (config)#mac vrf elan500 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.1:500 | Configuring Route-Distinguisher value 10.143.73.1:500 |
| (config-vrf)#route-target both 500:500 | Configuring import and export value as 500:500 |
| (config)# ip vrf vrf600 | Enter VRF mode |
| (config-vrf)# tunnel-select-policy PE1-PE3 | Configure tunnel-policy to be used for the VRF. |
| (config-vrf)# rd 10.143.73.1:600 | Configuring Route-Distinguisher value 10.143.73.1:600 |
| (config-vrf)#route-target both 600:600 | Configuring import and export value as 600:600 |
| (config-vrf)#exit | Exiting VRF Mode. |

PE1: EVPN and MAC-VRF mapping for ELAN and ELINE

| | |
|--|---|
| (config)#evpn mpls id 2 xconnect target-mpls-id 252 | Configure the EVPN-ELINE identifier with source identifier 2 and target identifier 252. |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier. |
| (config-evpn-mpls)#exit | Exiting from evpn-mpls mode |
| (config)# evpn mpls id 500 | Configure evpn mpls id 500 |
| (config-evpn-mpls)#tunnel-select-policy PE1-P2-PE3 | Map tunnel-policy to EVI |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp elan500 | Mapping vrf " elan500" to EVPN-VPLS identifier. |

PE1: Access side Configuration

| | |
|---|--|
| (config)#interface xe8 | Enter the Interface mode for xe8 |
| (config-if)# switchport | Configure interface as switch port |
| (config-if)# mpls-l2-circuit pe1-to-pe3 service-template pe1-to-pe3 primary | Bind the interface to the VC with service template |
| (config-if)# mpls-vpls vpls-pe1-to-pe3pe4-1 service-template vpls-pe1-to-pe3pe4-1 | Bind the VPLS instance to the interface |

| | |
|---|--|
| (config-if-vpls)# exit-if-vpls | Exit VPLS attachment-circuit mode |
| (config-if)# mpls-vpls vpls-pe1-to-pe3pe4-2 service-template vpls-pe1-to-pe3pe4-2 | Bind the VPLS instance to the interface |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if-vpls)# exit-if-vpls | Exit VPLS attachment-circuit mode |
| (config-if)# exit | Exit from the interface |
| (config)#interface xe8.2 switchport | Configure Interface xe8.2 as switchport |
| (config-if)#encapsulation dot1q 2 | Setting Encapsulation to dot1q with VLAN ID 2 |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if)#map vpn-id 2 | Map vpn-id 2 to interface xe8.2 (VPWS). |
| (config-access-if)#exit | Exiting out of access interface mode. |
| (config)#interface xe8.500 switchport | Configure Interface xe8.500 as switchport |
| (config-if)#encapsulation dot1q 500 | Setting Encapsulation to dot1q with VLAN ID 500 |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if)#map vpn-id 500 | Map vpn-id 500 to interface xe8.500 (VPLS). |
| (config-access-if)#exit | Exiting out of access interface mode. |
| (config)#interface xe8.600 | Enter Interface mode xe8.600 |
| (config-if)# ip vrf forwarding vrf600 | Bind the VRF instance to the interface |
| (config-if)# ip address 170.1.1.1/24 | Configure IP address |
| (config-if)# ipv6 address 1111::1/64 | Configure ipv6 address |
| (config-if)#encapsulation dot1q 600 | Setting Encapsulation to dot1q with VLAN ID 600 |
| (config-if)#exit | Exiting interface Mode. |
| (config)# router bgp 100 | Enter BGP router mode |
| (config-router)#address-family ipv4 vrf vrf600 | Configure VRF address family |
| (config-router-af) redistribute connected | Redistribute connected router |
| (config-router-af)#neighbor 170.1.1.2 remote-as 200 | Configure neighbor with remote-as |
| (config-router-af)#neighbor 170.1.1.2 activate | Activate the VPN neighbour |
| (config-router-af)exit-address-family | Exit VRF address family |
| (config-router)#address-family ipv6 vrf vrf600 | Configure ipv6 VRF address family |

| | |
|---|---|
| (config-router-af) redistribute connected | Redistribute connected router |
| (config-router-af) #neighbor 1111::2 remote-as 200 | Configure neighbor with remote-as |
| (config-router-af) #neighbor 1111::2 activate | Activate the VPN neighbour |
| (config-router-af) exit-address-family | Exit VRF address family |
| (config-router-af) exit | End |
| (config-router) #commit | Commit the candidate configuration to the running configuration |

P1: Loopback Interface

| | |
|---|--|
| (config) #interface lo | Enter the Interface mode for the loopback interface. |
| (config-if) #ip address 10.143.73.5/32 secondary | Configure IP address on loopback interface. |
| (config-if) #prefix-sid index 5 | Configure prefix sid index value |
| (config-if) #exit | Exit interface mode |

P1: Interface Configuration Network Side

| | |
|---|--|
| (config) #interface xe15 | Enter the Interface mode for xe15 |
| (config-if) # ip address 10.255.128.3/31 | Configure IP address on the interface. |
| (config-if) # mtu 9900 | Configure mtu |
| (config-if) # label-switching | Enable label switching on the interface. |
| (config-if) #exit | Exit interface mode |
| (config) #interface xe7 | Enter the Interface mode for xe7. |
| (config-if) # ip address 10.255.128.23/31 | Configure IP address on the interface. |
| (config-if) # label-switching | Enable label switching on the interface. |
| (config-if) #exit | Exit interface mode |

P1: OSPF Configuration

| | |
|--|---|
| (config) #router ospf 100 | Enter the Router OSPF mode. |
| (config-router) #ospf router-id 10.143.73.5 | Router-id configurations |
| (config-router) # bfd all-interfaces | Enable the OSPF enabled interfaces with bfd |
| (config-router) # network 10.143.73.5/32 area | |

| | |
|---|---|
| 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# network 10.255.128.2/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# network 10.255.128.22/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# ospf segment-routing global block 16000 17000 | Enable SRGB under ospf process 100 |
| (config-router)#segment-routing mpls | Enable segment routing |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit the candidate configuration to the running configuration |

P2: Loopback Interface

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.6/32 secondary | Configure IP address on loopback interface. |
| (config-if)#prefix-sid index 6 | Configure prefix sid index value |
| (config-if)#exit | Exit interface mode |

P2: Interface Configuration Network Side

| | |
|--|--|
| (config)#interface xe4 | Enter the Interface mode for xe4 |
| (config-if)# ip address 10.255.128.11/31 | Configure IP address on the interface. |
| (config-if)# mtu 9900 | Configure mtu |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter the Interface mode for xe5. |
| (config-if)# ip address 10.255.128.25/31 | Configure IP address on the interface. |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |

P2: OSPF Configuration

| | |
|---|-----------------------------|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.6 | Router-id configurations |

| | |
|---|---|
| (config-router)# bfd all-interfaces | Enable the OSPF enabled interfaces with bfd |
| (config-router)# network 10.143.73.6/32 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# network 10.255.128.10/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# network 10.255.128.24/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# ospf segment-routing global block 16000 17000 | Enable SRGB under ospf process 100 |
| (config-router)#segment-routing mpls | Enable segment routing |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit the candidate configuration to the running configuration |

PE3: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)# ip address 10.143.73.3/32 secondary | Configure IP address on loopback interface. |
| (config-if)#prefix-sid index 3 | Configure prefix sid index value |
| (config-if)#exit | Exit interface mode |

PE3: Global EVPN MPLS Command

| | |
|--|---|
| #configure terminal | Enter configuration mode. |
| (config)#evpn mpls enable | Enable EVPN MPLS. Note: Reload is required after Enabling/ Disabling EVPN MPLS Feature. |
| (config)#evpn mpls vtep-ip-global 10.143.73.3 | Configuring vtep global ip to loopback IP. |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | End configuration mode |
| #write memory | Saving configs to startup config |
| #reload | Reload device after enabling evpn |

PE3: LDP Configuration

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.3 | Set the router ID to IP address 10.143.73.3. |
| (config-router)#transport-address ipv4 10.143.73.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)# targeted-peer ipv4 10.143.73.1 | Configure targeted peer |
| (config-router-targeted-peer)# exit-targeted-peer-mode | Exit from router target peer and LDP mode |

PE3: Interface Configuration Network Side

| | |
|--|--|
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)# ip address 10.255.128.24/31 | Configure IP address on the interface. |
| (config-if)#mtu 9900 | Configure mtu at interface level |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)# enable-ldp ipv4 | Enable ldp on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe7 | Enter the Interface mode for xe7. |
| (config-if)# ip address 10.255.128.22/31 | Configure IP address on the interface. |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)# enable-ldp ipv4 | Enable ldp on the interface. |
| (config-if)#exit | Exit interface mode |

PE3: OSPF Configuration

| | |
|--|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.3 | Router-id configurations |
| (config-router)# bfd all-interfaces | Enable the OSPF enabled interfaces with bfd |
| (config-router)# network 10.143.73.3/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)# network 10.255.128.22/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# network 10.255.128.24/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# ospf segment-routing | Enable SRGB under ospf process 100 |

| | |
|--------------------------------------|---|
| global block 16000 17000 | |
| (config-router)#segment-routing mpls | Enable segment routing |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |

PE3: BGP Configuration

| | |
|--|--|
| (config)# router bgp 100 | Enter the Router BGP mode, ASN: 100. |
| (config-router)# neighbor 10.143.73.1 remote-as 100 | Configuring PE3 as iBGP neighbor using it's loopback ip. |
| (config-router)# neighbor 10.143.73.1 fall-over bfd multihop | Configure neighbor bfd multihop |
| (config-router)# neighbor 10.143.73.1 update-source lo | Source of routing updates as loopback |
| (config-router)# neighbor 10.143.73.1 advertisement-interval 0 | Configure advertisement interval |
| (config-router)# address-family vpnv4 unicast | Enter VPNv4 Address family mode |
| (config-router-af)# neighbor 10.143.73.1 activate | Enabling VPNv4 Address family for neighbor.. |
| (config-router-af)# exit-address-family | Exit Address-family mode. |
| (config-router)# address-family l2vpn vpls | Enter l2vpn Address family mode |
| (config-router-af)# neighbor 10.143.73.1 activate | Enabling l2vpn Address family for neighbor.. |
| (config-router-af)# exit-address-family | Exit Address-family mode. |
| (config-router)# address-family l2vpn evpn | Enter evpn Address family mode |
| (config-router-af)# neighbor 10.143.73.1 activate | Enabling evpn Address family for neighbor.. |
| (config-router-af)# exit-address-family | Exit Address-family mode. |
| (config-router)# address-family vpnv6 unicast | Enter VPNv6 Address family mode |
| (config-router-af)# neighbor 10.143.73.1 activate | Enabling VPNv6 Address family for neighbor.. |
| (config-router-af)# exit-address-family | Exit Address-family mode. |

PE3: VPWS Configuration

| | |
|---|---|
| (config)# mpls l2-circuit pe3-to-pe1 2002 10.143.73.1 | Configure the VC for PE-1. In this example, pe3-to-pe1 is the VC name, 2002 is the VC ID, and |
|---|---|

| | |
|--------------------------------------|---|
| | 10.143.73.1 is the VC endpoint IP address. |
| (config)#exit | Exiting from config mode |
| (config)#service-template pe3-to-pe1 | Template configuration. |
| (config-svc)# match outer-vlan 2002 | Match criteria under template configuration |
| (config)#exit | Exiting from config mode |

PE3: VPLS Configuration

| | |
|---|--|
| (config)# mpls vpls vpls-pe3-to-pe1pe4-1 2502 | Enter VPLS config mode |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type Ethernet | Type ethernet configuration for VPLS |
| (config-vpls-sig)# vpls-peer 10.143.73.1 | Configure VPLS Peer |
| (config-vpls-sig)#exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#service-template vpls-pe3-to-pe1pe4-1 | Template configuration. |
| (config-svc)# match outer-vlan 2502 | Match criteria under template configuration |
| (config-svc)# rewrite ingress translate 2600 outgoing-tpid dot1.q | Action performed for service template. |
| (config)# mpls vpls vpls-pe3-to-pe1pe4-2 2503 | Enter VPLS config mode |
| (config-vpls)#signaling bgp | Enter the Signaling bgp mode for BGP VPLS. |
| (config-vpls-sig)# ve-id 4 | Configure VE ID, which is mandatory for BGP VPLS, otherwise, signaling does not take place. VE ID should be unique per VPLS instance |
| (config-vpls-sig)#exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#service-template vpls-pe3-to-pe1pe4-2 | Template configuration. |
| (config-svc)# match double-tag outer-vlan 2503 inner-vlan 2504 | Match criteria under template configuration |
| (config-svc)# rewrite ingress push 2505 | Action performed for service template. |

PE3: MAC and IP VRF Configuration

| | |
|-------------------------------|---|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.3:2 | Configuring Route-Distinguisher value 10.143.73.3:2 |

| | |
|--|---|
| (config-vrf)#route-target both 2:2 | Configuring import and export value as 2:2 |
| (config-vrf)#exit | Exiting VRF Mode. |
| (config)#mac vrf elan500 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.3:500 | Configuring Route-Distinguisher value 10.143.73.3:500 |
| (config-vrf)#route-target both 500:500 | Configuring import and export value as 500:500 |
| (config)# ip vrf vrf600 | Enter VRF mode |
| (config-vrf)# rd 10.143.73.3:600 | Configuring Route-Distinguisher value 10.143.73.3:600 |
| (config-vrf)#route-target both 600:600 | Configuring import and export value as 600:600 |
| (config-vrf)#exit | Exiting VRF Mode. |

PE3: EVPN and mac vrf mapping Configuration

| | |
|--|---|
| (config)#evpn mpls id 252 xconnect target-mpls-id 2 | Configure the EVPN-ELINE identifier with source identifier 252 and target identifier 2. |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier. |
| (config-evpn-mpls)#exit | Exiting from evpn-mpls mode |
| (config)# evpn mpls id 500 | Configure evpn mpls id 500 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp elan500 | Mapping vrf " elan500" to EVPN-VPLS identifier. |

PE3: Access side Configuration

| | |
|---|---|
| (config)#interface xe4 | Enter the Interface mode for xe8 |
| (config-if)# switchport | Configure interface as switch port |
| (config-if)# mpls-l2-circuit pe3-to-pe1 service-template pe3-to-pe1 primary | Bind the VPWS instance to the interface |
| (config-if)# mpls-vpls vpls-pe3-to-pelpe4-1 service-template vpls-pe3-to-pelpe4-1 | Bind the VPLS instance to the interface |
| (config-if-vpls)# exit-if-vpls | Exit VPLS attachment-circuit mode |
| (config-if)# mpls-vpls vpls-pe3-to-pelpe4-2 service-template vpls-pe3-to-pelpe4-2 | Bind the VPLS instance to the interface |
| (config-if-vpls)# exit-if-vpls | Exit VPLS attachment-circuit mode |
| (config-if)# exit | Exit from the interface |
| (config)#interface xe4.2 switchport | Configure Interface xe4.2 as switchport |

| | |
|---|---|
| (config-if)#encapsulation dot1q 2 | Setting Encapsulation to dot1q with VLAN ID 2 |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if)#map vpn-id 252 | Map vpn-id 252 to interface xe4.2 (VPWS). |
| (config-access-if)#exit | Exiting out of access interface mode. |
| (config)#interface xe4.500 switchport | Configure Interface xe4.500 as switchport |
| (config-if)#encapsulation dot1q 500 | Setting Encapsulation to dot1q with VLAN ID 500 |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if)#map vpn-id 500 | Map vpn-id 500 to interface xe4.500 (VPLS). |
| (config-access-if)#exit | Exiting out of access interface mode. |
| (config)#interface xe4.600 | Enter Interface mode xe4.600 |
| (config-if)# ip vrf forwarding vrf600 | Bind the VRF instance to the interface |
| (config-if)# ip address 180.1.1.1/24 | Configure IP address |
| (config-if)# ipv6 address 2222::1/64 | Configure ipv6 address |
| (config-if)#encapsulation dot1q 600 | Setting Encapsulation to dot1q with VLAN ID 600 |
| (config-if)#exit | Exiting interface Mode. |
| (config)# router bgp 100 | Enter BGP router mode |
| (config-router)#address-family ipv4 vrf vrf600 | Configure VRF address family |
| (config-router-af) redistribute connected | Redistribute connected router |
| (config-router-af)#neighbor 180.1.1.2 remote-as 300 | Configure neighbor with remote-as |
| (config-router-af)#neighbor 180.1.1.2 activate | Activate the VPN neighbour |
| (config-router-af) exit-address-family | Exit VRF address family |
| (config-router)#address-family ipv6 vrf vrf600 | Configure ipv6 VRF address family |
| (config-router-af) redistribute connected | Redistribute connected router |
| (config-router-af)#neighbor 2222::2 remote-as 300 | Configure neighbor with remote-as |
| (config-router-af)#neighbor 2222::2 activate | Activate the VPN neighbour |
| (config-router-af) exit-address-family | Exit VRF address family |
| (config-router-af) end | End |
| (config)#commit | Commit the candidate configuration to the running configuration |

Validation

CLI's : Show segment-routing policy, show tunnel-policy, show tunnel-policy service details, show mpls forwarding-table

PE1

```
PE1#show segment-routing policy
```

| Policy-Name | Color | End-point | State | Forwarding-Info |
|-------------|-------|-------------|-------|-----------------|
| PE1-P1-PE3 | 1 | 10.143.73.3 | UP | 3/16003/xe16 |
| PE1-P2-PE3 | 2 | 10.143.73.3 | UP | 3/16003/xe4 |

VPWS

```
PE1#show mpls vc-table
```

| VC-ID | Vlan-ID | Inner-Vlan-ID | Access-Intf | Network-Intf | Out Label | Tunnel-Label | Nexthop | Status |
|-------|---------|---------------|-------------|--------------|-----------|--------------|---------|--------|
| 2002 | N/A | N/A | xe8 | xe16 | 25600 | 3 | | Active |

```
PE1#show mpls l2-circuit
```

```
MPLS Layer-2 Virtual Circuit: pe1-to-pe3, id: 2002 PW-INDEX: 2 service-tpid: dot1.q
Tunnel-Policy: PE1-P1-PE3 >>> VPWS is up by using Applied Tunnel Policy
Endpoint: 10.143.73.3
Control Word: 0
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe8
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : pe1-to-pe3
Match criteria : 2002
```

VPLS

```
PE1#show mpls vpls mesh
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | |
|---------|-------------|--------------|----------|--------------|-----------|---------|----------|---|
| 2502 | 10.143.73.3 | 3 | 26241 | xe4 | 25601 | 2/Up | 1 | L |
| DP | Active | | | | | | | |
| 2503 | 10.143.73.3 | 3 | 25604 | xe16 | 24962 | 2/Up | 3 | B |
| GP | Active | | | | | | | |

```
PE1#show mpls vpls detail
```

```
Virtual Private LAN Service Instance: vpls-pe1-to-pe3pe4-1, ID: 2502
SIG-Protocol: LDP
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Configured interfaces:
Interface: xe8
Service-template : vpls-pe1-to-pe3pe4-1
Match criteria : 2502
Action type : Translate
Action value : 2600
```

```

Outgoing tpid : dot1.q

Mesh Peers:
  10.143.73.3 (Up)
  Tunnel-Policy: PE1-P2-PE3 >>> VPLS-LDP is up using applied tunnel policy

Virtual Private LAN Service Instance: vpls-pe1-to-pe3pe4-2, ID: 2503
SIG-Protocol: BGP
  Route-Distinguisher :100:2503
  Route-Target :100:2503
  VE-ID :3
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Configured interfaces:
  Interface: xe8
Service-template : vpls-pe1-to-pe3pe4-2
Match criteria : 2503/2504
Action type : Push
Action value : 2505

Mesh Peers:
  10.143.73.3 (Up)
  Tunnel-Policy: PE1-P1-PE3 VPLS-BGP is up using applied tunnel policy

```

EVPN E-LAN

```

PE1#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels
(*) in Policy - tunnel-policy inherited from mac-vrf
Destination      Status      evpn-id  Policy      Network-Intf Tunnel-Label MC-
  Local          Remote
Label UC-Label MC-Label UC-Label
=====
10.143.73.3      Installed  500      PE1-P2-
PE3      xe4          3          640      17          640      --
Total number of entries are 1

Here Tunnel Policy is applied under VNID level.

```

EVPN E-LINE

```

PE1#show evpn mpls xconnect tunnel label
EVPN-MPLS Network tunnel labels
(*) in Policy - tunnel-policy inherited from mac-vrf
              Local      Remote
Destination   Status   VPWS-ID  VPWS-ID  Policy      Network-Intf Tunnel-
  Local        Remote
Label MC-Label UC-Label MC-Label UC-Label
=====
10.143.73.3    Installed  2          252      PE1-P1-PE3(*) xe16          3          --
  16          --          16
Total number of entries are 1

Here Tunnel Policy is applied under mac-vrf level.

```

L3VPN and 6VPE

```

PE1#show mpls vrf-table
Output for IPv4 VRF table with id: 4
Primary FTN entry with FEC: 180.1.1.0/24, id: 1, row status: Active, Tunnel-Policy: PE1-P1-PE3
Owner: BGP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Transport Tunnel id: 1, Protected LSP id: 0, QoS Resource id: 0, Description: N/A, Color: 1
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 8
Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 8, owner: BGP, Stale: NO, BGP out intf: xe16, transport out intf: xe16,
out label: 25024
Nexthop addr: 10.143.73.3          cross connect ix: 6, op code: Push and Lookup

Output for IPv6 VRF table with id: 4
Primary FTN entry with FEC: 2222::/64, id: 13, row status: Active, Tunnel-Policy: PE1-P1-PE3
Owner: BGP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Transport Tunnel id: 1, Protected LSP id: 0, QoS Resource id: 0, Description: N/A, Color: 1
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 8
Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 8, owner: BGP, Stale: NO, BGP out intf: xe16, transport out intf: xe16,
out label: 25024
Nexthop addr: 10.143.73.3          cross connect ix: 6, op code: Push and Lookup

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code   FEC           FTN-ID   Nhlfe-ID   Tunnel-id   Pri   LSP-Type   Out-Label   Out-
Intf   ELC   Nexthop
P>    10.143.73.3/32   4         10         1         Yes   LSP_
DEFAULT 3         xe16       No        10.255.128.3
P>    10.143.73.3/32   5         11         2         Yes   LSP_
DEFAULT 3         xe4        No        10.255.128.11
O     10.143.73.3/32   1         3          0         Yes   LSP_
DEFAULT 16003        xe16       No        10.255.128.3
O>    10.143.73.5/32   2         5          0         Yes   LSP_
DEFAULT 3         xe16       No        10.255.128.3
O>    10.143.73.6/32   3         6          0         Yes   LSP_
DEFAULT 3         xe4        No        10.255.128.11

PE1#show mpls forwarding-table detail
FEC prefix: 10.143.73.3/32, FTN-ID: 4
Owner: SR_POLICY, FTN type: Regular, State: Installed
Tunnel-Name: N/A, Tunnel-id: 1, Color: 1
LSP-ID: N/A, LSP-type: Primary
NHLFE-id: 10
Out-Label: 3, Out-Intf: xe16, Nexthop: 10.255.128.3
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

FEC prefix: 10.143.73.3/32, FTN-ID: 5
Owner: SR_POLICY, FTN type: Regular, State: Installed
Tunnel-Name: N/A, Tunnel-id: 2, Color: 2
LSP-ID: N/A, LSP-type: Primary
NHLFE-id: 11
Out-Label: 3, Out-Intf: xe4, Nexthop: 10.255.128.11
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

FEC prefix: 10.143.73.3/32, FTN-ID: 1
Owner: OSPF-SR, FTN type: Regular, State: Not Selected
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
NHLFE-id: 3

```

```

Out-Label: 16003, Out-Intf: xe16, Nexthop: 10.255.128.3
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

FEC prefix: 10.143.73.5/32, FTN-ID: 2
Owner: OSPF-SR, FTN type: Regular, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
NHLFE-id: 5
Out-Label: 3, Out-Intf: xe16, Nexthop: 10.255.128.3
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

FEC prefix: 10.143.73.6/32, FTN-ID: 3
Owner: OSPF-SR, FTN type: Regular, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
NHLFE-id: 6
Out-Label: 3, Out-Intf: xe4, Nexthop: 10.255.128.11
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No

```

Tunnel policy Output after applying under all services

```

PE1#show tunnel-policy
tunnel-policy PE1-P1-PE3
  color 1
  tunnel-mode dedicated
  tunnel-type any
Total Services attached:
  VRF      : 1
  VPWS     : 1
  VPLS     : 1
  EVPN     : 1

tunnel-policy PE1-P2-PE3
  color 2
  tunnel-mode dedicated
  tunnel-type any
Total Services attached:
  VRF      : 0
  VPWS     : 0
  VPLS     : 1
  EVPN     : 1

PE1#show tunnel-policy service details
List of services configured with tunnel-policy PE1-P1-PE3:
VPWS
  VPWS-Id: 2002, Name: pel-to-pe3 Status: UP, Mapped Tunnel: SR_POLICY, ftn_ix 4, color 1
L3VPN
  VRF id 4
    FTN id: 1, Prefix: 180.1.1.0/24, Status: Up, Label: 25024, Mapped Tunnel: SR_POLICY, ftn_ix 4,
color 1
    FTN id: 13, Prefix: 2222::/64, Status: Up, Label: 25024, Mapped Tunnel: SR_POLICY, ftn_ix 4,
color 1
VPLS
  VPLS-Id: 2503, Name: vpls-pel-to-pe3pe4-2, Status: UP, Mapped Tunnel: SR_POLICY, ftn_ix 4, color 1
EVPN
  EVPN-Id: 2, Tunnel count: 1
    Peer: 10.143.73.3, Mapped Tunnel: SR_POLICY, ftn_ix 4, color 1
List of services configured with tunnel-policy PE1-P2-PE3:
VPLS
  VPLS-Id: 2502, Name: vpls-pel-to-pe3pe4-1, Status: UP, Mapped Tunnel: SR_POLICY, ftn_ix 5, color 2
EVPN
  EVPN-Id: 500, Tunnel count: 1
    Peer: 10.143.73.3, Mapped Tunnel: SR_POLICY, ftn_ix 5, color 2

```

PE2 Validation

```
PE2#show mpls vc-table
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VC-ID Vlan-ID Inner-Vlan-ID Access-Intf Network-Intf Out Label Tunnel-Label Nexthop Status
2002 N/A N/A xe22 xe7 26880 16001 10.143.73.1 Active

PE2#show mpls l2-circuit
MPLS Layer-2 Virtual Circuit: pe3-to-pe1, id: 2002 PW-INDEX: 2 service-tpid: dot1.q
Endpoint: 10.143.73.1
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe4
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : pe3-to-pe1
Match criteria : 2002

PE2#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID Peer Addr Tunnel-Label In-Label Network-Intf Out-Label Lkps/St PW-INDEX SIG-Protocol Status
2502 10.143.73.1 16001 26881 xe7 26881 2/Up 1 LDP Active
2503 10.143.73.1 16001 25602 xe7 25603 2/Up 3 BGP Active

PE2#show mpls vpls detail
Virtual Private LAN Service Instance: vpls-pe3-to-pelpe4-1, ID: 2502
SIG-Protocol: LDP
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Configured interfaces:
Interface: xe4
Service-template : vpls-pe3-to-pelpe4-1
Match criteria : 2502
Action type : Translate
Action value : 2600
Outgoing tpid : dot1.q

Mesh Peers:
10.143.73.1 (Up)
Virtual Private LAN Service Instance: vpls-pe3-to-pelpe4-2, ID: 2503

SIG-Protocol: BGP
Route-Distinguisher :100:2503
Route-Target :100:2503
VE-ID :4
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
```

Configured interfaces:

Interface: xe4
 Service-template : vpls-pe3-to-pelpe4-2
 Match criteria : 2503/2504
 Action type : Push
 Action value : 2505

Mesh Peers:

10.143.73.1 (Up)

PE2#show evpn mpls tunnel label

EVPN-MPLS Network tunnel labels

(*) in Policy - tunnel-policy inherited from mac-vrf

(e) - Service mapped over MPLS Multipath/ECMP

```
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
```

Local Remote MPLS-Multipath Underlay

Destination Status VPN-ID Policy MC-Label UC-Label MC-Label UC-Label Grp-Name NHLFE-ix NW-Intf NW-Label

```
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
```

10.143.73.1 Installed 500 -- 640 17 640 -- -- -- xe7 16001

Total number of entries are 1

PE2#show evpn mpls xconnect tunnel label

EVPN-MPLS Network tunnel labels

(*) in Policy - tunnel-policy inherited from mac-vrf

(e) - Service mapped over MPLS Multipath/ECMP

```
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
```

Local Remote Local Remote MPLS-Multipath Underlay

Destination Status VPWS-ID VPWS-ID Policy UC-Label UC-Label Grp-Name NHLFE-ix NW-Intf NW-Label

```
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
```

10.143.73.1 Installed 252 2 -- 16 16 -- -- -- xe7 16001

Total number of entries are 1

PE2#show mpls vrf-table

Output for IPv4 VRF table with id: 4

Primary FTN entry with FEC: 170.1.1.0/24, id: 1, row status: Active, Tunnel-Policy: N/A

Owner: BGP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none

VRF id 4, BGP peer 10.143.73.1 BGP prefix 170.1.1.0

Transport Tunnel id: 0, Protected LSP id: 0, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 5

Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 5, owner: BGP, Stale: NO, BGP out intf: xe7, transport out intf: xe7, out label: 25664

Nexthop addr: 10.143.73.1 cross connect ix: 4, op code: Push and Lookup

Output for IPv6 VRF table with id: 4

Primary FTN entry with FEC: 1111::/64, id: 2, row status: Active, Tunnel-Policy: N/A

Owner: BGP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none

VRF id 4, BGP peer a8f:4901:: BGP prefix 1111::

Transport Tunnel id: 0, Protected LSP id: 0, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 5

Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 5, owner: BGP, Stale: NO, BGP out intf: xe7, transport out intf: xe7, out label: 25664

Nexthop addr: 10.143.73.1 cross connect ix: 4, op code: Push and Lookup

PE2#sh mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup

B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,

L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,

U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

Code FEC FTN-ID Nhlfe-ID Tunnel-id Pri LSP-Type Out-Label Out-Intf ELC Nexthop

O> 10.143.73.1/32 1 9 0 Yes LSP_DEFAULT 16001 xe7 No 10.255.128.23

```
O> 10.143.73.5/32 2 10 0 Yes LSP_DEFAULT 3 xe7 No 10.255.128.23
O> 10.143.73.6/32 3 14 0 Yes LSP_DEFAULT 3 xe5 No 10.255.128.25
PE2#show mpls forwarding-table detail
FEC prefix: 10.143.73.1/32, FTN-ID: 1
Owner: OSPF-SR, FTN type: Regular, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
NHLFE-id: 9
Out-Label: 16001, Out-Intf: xe7, Nexthop: 10.255.128.23
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No
FEC prefix: 10.143.73.5/32, FTN-ID: 2
Owner: OSPF-SR, FTN type: Regular, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
NHLFE-id: 10
Out-Label: 3, Out-Intf: xe7, Nexthop: 10.255.128.23
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No
FEC prefix: 10.143.73.6/32, FTN-ID: 3
Owner: OSPF-SR, FTN type: Regular, State: Installed
Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
LSP-ID: N/A, LSP-type: Primary
NHLFE-id: 14
Out-Label: 3, Out-Intf: xe5, Nexthop: 10.255.128.25
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0

Note:The following CLI will be used to change from dedicated mode to best effort
tunnel-policy PE1-P1-PE3
color 1
mode best-effort
```



Note: The following CLI will be used to change from dedicated mode to best effort
tunnel-policy PE1-P1-PE3
color 1
mode best-effort

Mapping RSVP Tunnel Name to L2VPN Service

This chapter shows configurations of mapping of rsvp tunnel-name to L2VPN service.

An MPLS Layer 2 Virtual Circuit (VC) is a point-to-point Layer 2 connection transported via MPLS on the service provider's network. The Layer 2 circuit is transported over a single Label Switched Path (LSP) tunnel between two Provider Edge (PE) routers

Virtual Private LAN Service (VPLS) is a way to provide Ethernet-based multipoint-to-multipoint communication over IP- MPLS networks. It allows geographically-dispersed sites to share an Ethernet broadcast domain by connecting sites through pseudowires. A set of Martini circuits is grouped by a common VPLS identifier to achieve this service objective.




Note: Starting with OcNOS version 7.0.0, sub-interfaces replace Service Templates as the preferred configuration method. While legacy Service Template configurations remain functional, the associated commands are now hidden and will be completely removed in a future release. To ensure future-proofing and continued manageability, please migrate all existing Service Templates to sub-interfaces. For more details, see [Appendix A: Service Template Configuration \(page 2631\)](#).

Topology

This topology will be applicable for both VPWS and VPLS services.

Figure 44. Mapping of RSVP Tunnel-name to L2VPN services



Configuring the Virtual Circuit (VC)



Note: Loopback addresses being used should be advertised through OSPF, or should be statically routed.

1. Configure the IP address and OSPF for the PE1, P (Provider), and PE2 routers.
2. Configure MPLS and LDP on PE1, P, and PE2, and LDP targeted peer for the PE1 and PE2 routers. (If RSVP is used for configuring trunks, LDP must be configured on PE1 and PE2, and RSVP must be configured on PE1, P, and PE2).
3. Configure the VC with trunk-name.
4. Bind the customer interface to the VC.

Configure IP Address and OSPF on Routers

Configure the IP addresses and OSPF on the PE1, P, and PE2 routers.

PE1

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#ip address 10.1.1.1/24 | Set the IP address of the interface to 10.1.1.1/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.1.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

P

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 9.9.9.9/32 secondary | Set the IP address of the loopback interface to 9.9.9.9/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#ip address 10.1.1.2/24 | Set the IP address of the interface to 10.1.1.2/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe13 | Specify the interface (xe13) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe13. |
| (config-if)#ip address 20.1.1.1/24 | Set the IP address of the interface to 20.1.1.1/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the |

| | |
|---|---|
| | Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| <code>(config-router)#network 10.1.1.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#network 20.1.1.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#network 9.9.9.9/32 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| <code>(config-router)#commit</code> | Commit the transaction. |

PE2

| | |
|--|---|
| <code>#configure terminal</code> | Enter configure mode. |
| <code>(config)#interface lo</code> | Specify the loopback interface (lo0) to be configured. |
| <code>(config-if)#ip address 2.2.2.2/32 secondary</code> | Set the IP address of the loopback interface to 2.2.2.2/32. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe13</code> | Specify the interface (xe13) to be configured. |
| <code>(config-if)#label-switching</code> | Enable label switching on interface xe13. |
| <code>(config-if)#ip address 20.1.1.2/24</code> | Set the IP address of the interface to 20.1.1.2/24. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router ospf 100</code> | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| <code>(config-router)#network 20.1.1.0/24 area 0</code> | Define the interface on which OSPF runs, and associate the area ID (0) with the interface. |
| <code>(config-router)#network 2.2.2.2/32 area 0</code> | Define the interface on which OSPF runs, and associate the area ID (0) with the interface. |
| <code>(config-router)#commit</code> | Commit the transaction. |

Configure MPLS, RSVP, and LDP Targeted Peer on Routers

Configure MPLS and LDP on PE1, P, and PE2, and LDP targeted peers on PE1 and PE2.



Note: If RSVP is used for configuring trunks, LDP must be configured on PE1 and PE2, and RSVP must be configured on PE1, P, and PE2.

PE1

| | |
|----------------------------------|-----------------------|
| <code>#configure terminal</code> | Enter configure mode. |
|----------------------------------|-----------------------|

| | |
|--|--|
| (config)#router ldp | Enter the Router mode. |
| (config-router)#router-id 1.1.1.1 | Configure the router-id |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Specify the targeted LDP peer on PE1. |
| (config-router-targeted-peer)# exit | Exit the Router targeted peer mode. |
| (config-router)#exit | Exit the Router mode. |
| (config)#router rsvp | Enter RSVP configuration mode for the router. |
| (config-router)#exit | Exit configuration mode of the router. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe1. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#rsvp-trunk t2 | Configure RSVP trunk t2 |
| (config-trunk)#to 2.2.2.2 | Configure PE2 as the end of trunk |
| (config-trunk)#commit | Commit the transaction. |

P

| | |
|---------------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#router rsvp | Enter RSVP configuration mode for the router. |
| (config)#exit | Exit configuration mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe13 | Specify the interface (xe13) to be configured. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#rsvp-trunk t5 | Configure RSVP trunk t5 |
| (config-trunk)#to 2.2.2.2 | Configure PE2 as the end of trunk |
| (config-trunk)#exit | Exit configuration mode |
| (config)#rsvp-trunk t6 | Configure RSVP trunk t6 |
| (config-trunk)#to 1.1.1.1 | Configure PE2 as the end of trunk |
| (config-trunk)#commit | Commit the transaction. |

PE2

| | |
|---------------------|-----------------------|
| #configure terminal | Enter configure mode. |
|---------------------|-----------------------|

| | |
|--|--|
| (config)#router ldp | Enter the Router mode. |
| (config-router)#router-id 2.2.2.2 | Configure the router-id |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Specify the targeted LDP peer on PE2. |
| (config-router-targeted-peer)# exit | Exit the Router targeted peer mode. |
| (config-router)#exit | Exit the Router mode. |
| (config)#router rsvp | Enter RSVP configuration mode for the router. |
| (config-router)#exit | Exit configuration mode of the router. |
| (config)#interface xe13 | Specify the interface(xe13) to be configured. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe13. |
| (config-if)#enable-rsvp | Enable RSVP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#rsvp-trunk t3 | Configure RSVP trunk t3 |
| (config-trunk)#to 1.1.1.1 | Configure PE1 as the end of trunk |
| (config-trunk)#commit | Commit the transaction. |

Configure VC

Configure the VC. Each VC ID uniquely identifies the Layer-2 circuit among all the Layer-2 circuits.



Note: Both PE routers (endpoints) must be configured with the same VC-ID (100 in this example).

PE1

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#mpls l2-circuit t2 100 2.2.2.2 | Configure the VC for PE2. In this example, t2 is the VC name, 200 is the VC ID, and 2.2.2.2 is the VC endpoint IP address |
| (config-pseudowire)#tunnel-name t2 | Configure the RSVP Trunk name as t2 |
| (config-pseudowire)#exit | Exit pseudowire config mode. |
| (config)#mpls l2-circuit t3 300 2.2.2.2 mode raw | Configure the VC for PE2. In this example, t3 is the VC-name, 300 is the VC ID and 2.2.2.2 is the VC endpoint IP address |
| (config-pseudowire)#tunnel-name t2 | Configure RSVP Trunk name as t2 |
| (config-pseudowire)#commit | Commit the transaction. |

PE2

| | |
|---------------------|-----------------------|
| #configure terminal | Enter configure mode. |
|---------------------|-----------------------|

| | |
|---|--|
| <code>(config)#mpls l2-circuit t2 100 1.1.1.1</code> | Configure the VC for PE1. In this example, t2 is the VC name, 200 is the VC ID, and 1.1.1.1 is the VC endpoint IP address |
| <code>(config-pseudowire)#tunnel-name t3</code> | Configure RSVP Tunnel name as t3 |
| <code>(config-pseudowire)#exit</code> | Exit pseudowire config mode. |
| <code>(config)#mpls l2-circuit t2 100 1.1.1.1</code> | Configure the VC for PE1. In this example, t2 is the VC name, 200 is the VC ID, and 1.1.1.1 is the VC endpoint IP address using rsvp trunk name t2 |
| <code>(config-pseudowire)#tunnel-name t2</code> | Configure RSVP Tunnel name as t2 |
| <code>(config-pseudowire)#exit</code> | Exit pseudowire config mode. |
| <code>(config)#mpls l2-circuit t3 300 1.1.1.1 mode raw</code> | Configure the VC for PE1. In this example, t3 is the VC name, 300 is the VC ID, and 1.1.1.1 is the VC endpoint IP address |
| <code>(config-pseudowire)#tunnel-name t3</code> | Configure RSVP Tunnel name as t3 |
| <code>(config-pseudowire)#commit</code> | Commit the transaction. |

Bind Customer Interface to VC

The following configuration allows only VLAN 2 and 3 traffic.

PE1

| | |
|--|---|
| <code>(config)#service-template ST1</code> | Create a service template ST1 |
| <code>(config-svc)#match outer-vlan 2</code> | Allow VLAN 2 traffic on this VC. |
| <code>(config-svc)#exit</code> | Exit the service template mode |
| <code>(config)#service-template ST2</code> | Create a service template ST2 |
| <code>(config-svc)#match outer-vlan 3</code> | Allow VLAN 3 traffic on this VC. |
| <code>(config-svc)#exit</code> | Exit the service template mode |
| <code>(config)#interface xe15</code> | Specify the interface (xe15) to be configured. |
| <code>(config-if)#switchport</code> | Switch to Layer-2 mode. |
| <code>(config-if)#mpls-l2-circuit t2 service-template ST1</code> | Bind the interface to the VC with service template. |
| <code>(config-if)#mpls-l2-circuit t3 service-template ST2</code> | Bind the interface to the VC with service template. |
| <code>(config-if)#commit</code> | Commit the transaction. |

PE2

| | |
|--|----------------------------------|
| <code>(config)#service-template ST1</code> | Create a service template ST1 |
| <code>(config-svc)#match outer-vlan 2</code> | Allow VLAN 2 traffic on this VC. |
| <code>(config-svc)#exit</code> | Exit the service template mode |
| <code>(config)#service-template ST2</code> | Create a service template ST2 |

| | |
|---|---|
| (config-svc)#match outer-vlan 3 | Allow VLAN 3 traffic on this VC. |
| (config-svc)#exit | Exit the service template mode |
| (config)#interface xe12 | Specify the interface (xe12) to be configured. |
| (config-if)#switchport | Switch to Layer-2 mode. |
| (config-if)#mpls-l2-circuit t2 service-template ST1 | Bind the interface to the VC with service template. |
| (config-if)#mpls-l2-circuit t3 service-template ST2 | Bind the interface to the VC with service template. |
| (config-if)#commit | Commit the transaction. |

Validation

Use the show ldp mpls-l2-circuit (Control Plane) command, and the show mpls vc-table (Forwarding Plane) command, to display complete information about the Layer 2 VC.

If the VC State is UP in the output from the show ldp mpls-l2 circuit command, and the Status is Active in the output of the show mpls vc-table command, a ping from CE1 to CE2 should be successful.

Below are the sample output for VPWS service with Tunnel name:

PE1

```

PE1#show mpls vc-table
VC-ID      Vlan-ID  Inner-Vlan-ID  Access-Intf  Network-Intf  Out Label  Tunnel-
Label      Nexthop   Status
100        N/A      N/A           xe15         xe1           24320      24321      2.2.2.2
Active
300        N/A      N/A           xe15         xe1           24321      24321      2.2.2.2
Active
PE1#

PE1#show ldp mpls-l2-circuit
Transport   Client    VC        VC          Local      Remote     Destination
VC ID      Binding   State     Type
300        xe15      UP        Ethernet
100        xe15      UP        Ethernet VLAN
PE1#

PE1#show mpls l2-circuit
MPLS Layer-2 Virtual Circuit: t2, id: 100  PW-INDEX: 1 Tunnel-Name: t2
Endpoint: 2.2.2.2
Control Word: 0
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe15
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST1
Match criteria : 2

MPLS Layer-2 Virtual Circuit: t3, id: 300  PW-INDEX: 2 Tunnel-Name: t2
Operating mode: Raw
Endpoint: 2.2.2.2
Control Word: 0
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe15
Virtual Circuit Type: Ethernet
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST2

```

Match criteria : 3

```
PE1#show ldp mpls-l2-circuit detail
PW ID: 300, VC state is up
Access IF: xe5,up,AC state is up
Session IF: xe2, state is up
Destination: 2.2.2.2, Peer LDP Ident: 2.2.2.2
Local vctype: ethernet, remote vctype :ethernet
Local groupid: 0, remote groupid: 0
Local label: 26881, remote label: 53760
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: Not-Applicable Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
MPLS VC UpTime : 00:02:12
```

```
PW ID: 100, VC state is up
Access IF: xe5,up,AC state is up
Session IF: xe2, state is up
Destination: 2.2.2.2, Peer LDP Ident: 2.2.2.2
Local vctype: vlan, remote vctype :vlan
Local groupid: 0, remote groupid: 0
Local label: 26880, remote label: 53761
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: Not-Applicable Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
MPLS VC UpTime : 00:02:12
```

PE2

```
PE2#sh mpls vc-table
VC-ID Vlan-ID Inner-Vlan-ID Access-Intf Network-Intf Out Label Tunnel-Label Nexthop Status
100 N/A N/A xe12 xe3 26880 25600 1.1.1.1 Active
300 N/A N/A xe12 xe3 26881 25600 1.1.1.1 Active
PE2#sh ldp mpls-l2-circuit
Transport Client VC VC Local Remote Destination
VC ID Binding State Type VC Label VC Label Address
300 xe12 UP Ethernet 53760 26881 1.1.1.1
100 xe12 UP Ethernet VLAN 53761 26880 1.1.1.1
```

```
PE2#sh mpls l2-circuit
MPLS Layer-2 Virtual Circuit: t2, id: 100 PW-INDEX: 1 service-tpid: dot1.q
Tunnel-Name: t3
Endpoint: 1.1.1.1
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe12
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST1
Match criteria : 2
```

```
MPLS Layer-2 Virtual Circuit: t3, id: 300 PW-INDEX: 2 service-tpid: dot1.q
Tunnel-Name: t3
```

Operating mode: Raw


```
Endpoint: 1.1.1.1
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe12
Virtual Circuit Type: Ethernet
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST2
Match criteria : 3

PE2#show ldp mpls-l2-circuit detail
PW ID: 300, VC state is up
Access IF: xe12,up,AC state is up
Session IF: xe3, state is up
Destination: 1.1.1.1, Peer LDP Ident: 1.1.1.1
Local vctype: ethernet, remote vctype :ethernet
Local groupid: 0, remote groupid: 0
Local label: 53760, remote label: 26881
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: Not-Applicable Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
MPLS VC UpTime : 00:05:34

PW ID: 100, VC state is up
Access IF: xe12,up,AC state is up
Session IF: xe3, state is up
Destination: 1.1.1.1, Peer LDP Ident: 1.1.1.1
Local vctype: vlan, remote vctype :vlan
Local groupid: 0, remote groupid: 0
Local label: 53761, remote label: 26880
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: Not-Applicable Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
MPLS VC UpTime : 00:05:34

These additional commands can also be used to display information about the Layer 2 virtual circuits.
show ldp mpls-l2-circuit detail
show ldp mpls-l2-circuit VC-ID
show ldp mpls-l2-circuit VC-ID detail
show mpls l2-circuit
```

Configuring a MPLS Static Layer-2 VC

1. Configure the VC with the manual option using tunnel name
2. Configure the VC FIB entry.
3. Bind the VC; all steps are in the configurations that follow.

PE1

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| PE1(config)#mpls l2-circuit t5 500 2.2.2.2 | Configure the VC for PE1 |
| PE1(config-pseudowire)#tunnel-name t2 | Configure the RSVP Tunnel name as t2 |
| PE1(config-pseudowire)#manual-pseudowire | Configure the VC as manual (no signaling is used) |
| PE1(config-pseudowire)#exit | Exit pseudowire config mode. |
| PE1(config)#service-template ST5 | Create a service template ST5 |
| PE1(config-svc)#match outer-vlan 5 | Configure single match criteria vlan 5 |
| PE1(config-svc)#exit | Exit the service template mode |
| PE1(config)#interface xe15 | Access interface xe15 |
| (config-if)#switchport | Switch to Layer-2 mode. |
| PE1(config-if)#mpls-l2-circuit t5 service-template ST5 | Bind the interface to the VC with service template. |
| PE1(config-if)#exit | Exit interface mode |
| PE1(config-if)#commit | Commit the transaction. |
| PE1(config)#mpls l2-circuit-fib-entry 500 1000 2000 2.2.2.2 xe1 xe15 | Add an FTN entry; where 1000 is the incoming label, 2000 is the outgoing label, 2.2.2.2 is the endpoint, xe1 is the Provider facing interface name, and xe15 is access interface name |
| PE1(config)#commit | Commit the transaction. |

PE2

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| PE2(config)#mpls l2-circuit t5 500 1.1.1.1 | Configure the VC for PE2 |
| PE2(config-pseudowire)#tunnel-name t3 | Configure RSVP Tunnel name as t3 |
| PE2(config-pseudowire)#manual-pseudowire | Configure VC as manual (no signaling used) |
| PE2(config-pseudowire)#exit | Exit pseudowire config mode. |
| PE2(config)#service-template ST5 | Create a service template ST5 |
| PE2(config-svc)#match outer-vlan 5 | Configure single match criteria vlan 5 |
| PE2(config-svc)#exit | Exit the service template mode |
| PE2(config)#interface xe12 | Access interface xe12 |
| (config-if)#switchport | Switch to Layer-2 mode. |
| PE2(config-if)#mpls-l2-circuit t5 service-template ST5 | Bind the interface to the VC with service template. |
| PE2(config-if)#exit | Exit interface mode. |
| PE2(config-if)#commit | Commit the transaction. |
| PE1(config)#mpls l2-circuit-fib-entry 500 2000 1000 1.1.1.1 xe13 xe12 | Add an FTN entry; where 2000 is the incoming label, 1000 is the outgoing label, 1.1.1.1 is the endpoint, |

| | |
|--------------------|--|
| | xe12 is the Provider facing interface name, and xe13 access interface name |
| PE1(config)#commit | Commit the transaction. |
| PE2(config)#exit | Exit configure mode |

Validation

This example shows number of configured VCs and its status.

PE1

```

PE1#show mpls vc-table
VC-ID      Vlan-ID  Inner-Vlan-ID  Access-Intf  Network-Intf  Out Label  Tunnel-
Label  Nexthop      Status
100      N/A        N/A           xe15         xe1           24320      24321      2.2.2.2
  Active
300      N/A        N/A           xe15         xe1           24321      24321      2.2.2.2
  Active
500      N/A        N/A           xe15         xe1           2000       24321      2.2.2.2
  Active

PE1#show mpls l2-circuit
MPLS Layer-2 Virtual Circuit: t2, id: 100  PW-INDEX: 1 Tunnel-Name: t2
Endpoint: 2.2.2.2
Control Word: 0
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe15
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST1
Match criteria : 2

MPLS Layer-2 Virtual Circuit: t3, id: 300  PW-INDEX: 2 Tunnel-Name: t2
Operating mode: Raw
Endpoint: 2.2.2.2
Control Word: 0
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe15
Virtual Circuit Type: Ethernet
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST2
Match criteria : 3

MPLS Layer-2 Virtual Circuit: t5, id: 500  PW-INDEX: 3 Tunnel-Name: t2
Endpoint: 2.2.2.2
Control Word: 0
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe15
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST5
Match criteria : 5

PE1#show ldp mpls-l2-circuit detail
PW ID: 300, VC state is up

```

```

Access IF: xe15,up,AC state is up
Session IF: xe1, state is up
Destination: 2.2.2.2, Peer LDP Ident: 2.2.2.2
Local vctype: ethernet, remote vctype :ethernet
Local groupid: 0, remote groupid: 0
Local label: 26881, remote label: 53760
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: Not-Applicable Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
MPLS VC UpTime : 00:15:48
PW ID: 100, VC state is up
Access IF: xe5,up,AC state is up
Session IF: xe1, state is up
Destination: 2.2.2.2, Peer LDP Ident: 2.2.2.2
Local vctype: vlan, remote vctype :vlan
Local groupid: 0, remote groupid: 0
Local label: 26880, remote label: 53761
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: Not-Applicable Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
MPLS VC UpTime : 00:15:48

```

PE2

```

OcNOS#sh mpls vc-table

VC-ID Vlan-ID Inner-Vlan-ID Access-Intf Network-Intf Out Label Tunnel-Label Nexthop Status
100 N/A N/A xe12 xe3 26880 25600 1.1.1.1 Active
300 N/A N/A xe12 xe3 26881 25600 1.1.1.1 Active
500 N/A N/A xe12 xe3 1000 25600 1.1.1.1 Active
PE2#sh mpls l2-circuit

MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe12
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST1
Match criteria : 2

MPLS Layer-2 Virtual Circuit: t3, id: 300 PW-INDEX: 2 service-tpid: dot1.q
Tunnel-Name: t3
Operating mode: Raw
Endpoint: 1.1.1.1
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe12
Virtual Circuit Type: Ethernet
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST2
Match criteria : 3

MPLS Layer-2 Virtual Circuit: t5, id: 500 PW-INDEX: 3 service-tpid: dot1.q
Tunnel-Name: t3
Endpoint: 1.1.1.1
Control Word: 0

```

```
Flow Label Status: Disabled, Direction: None, Static: No
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe12
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST5
Match criteria : 5

OcNOS#show ldp mpls-l2-circuit detail
PW ID: 300, VC state is up
Access IF: xe12,up,AC state is up
Session IF: xe3, state is up
Destination: 1.1.1.1, Peer LDP Ident: 1.1.1.1
Local vctype: ethernet, remote vctype :ethernet
Local groupid: 0, remote groupid: 0
Local label: 53760, remote label: 26881
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: Not-Applicable Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled

MPLS Layer-2 Virtual Circuit: t2, id: 100 PW-INDEX: 1 service-tpid: dot1.q
Tunnel-Name: t3
Endpoint: 1.1.1.1
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No

MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe12
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST1
Match criteria : 2

MPLS Layer-2 Virtual Circuit: t3, id: 300 PW-INDEX: 2 service-tpid: dot1.q
Tunnel-Name: t3
Operating mode: Raw
Endpoint: 1.1.1.1
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe12
Virtual Circuit Type: Ethernet
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST2

Match criteria : 3
MPLS Layer-2 Virtual Circuit: t5, id: 500 PW-INDEX: 3 service-tpid: dot1.q
Tunnel-Name: t3
Endpoint: 1.1.1.1
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe12
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : ST5
```

```

Match criteria : 5
OcNOS#show ldp mpls-l2-circuit detail
PW ID: 300, VC state is up
Access IF: xe12,up,AC state is up
Session IF: xe3, state is up
Destination: 1.1.1.1, Peer LDP Ident: 1.1.1.1
Local vctype: ethernet, remote vctype :ethernet
Local groupid: 0, remote groupid: 0
Local label: 53760, remote label: 26881
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: Not-Applicable Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
MPLS VC UpTime : 00:17:41

```

```

PW ID: 100, VC state is up
Access IF: xe12,up,AC state is up
Session IF: xe3, state is up
Destination: 1.1.1.1, Peer LDP Ident: 1.1.1.1
Local vctype: vlan, remote vctype :vlan
Local groupid: 0, remote groupid: 0
Local label: 53761, remote label: 26880
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: Not-Applicable Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
MPLS VC UpTime : 00:17:41

```

These additional commands can also be used to display information about the Layer 2 virtual circuits.

```

show ldp mpls-l2-circuit detail
show ldp mpls-l2-circuit VC-ID
show ldp mpls-l2-circuit VC-ID detail
show mpls l2-circuit

```

Configure Dynamic VPLS

PE1: LDP VPLS Configuration

| | |
|--|--|
| (config)#mpls vpls v1 25 | Enter VPLS config mode |
| (config-vpls)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 2.2.2.2 tunnel-name t2 | Configure VPLS Peer with trunk-name t2 |
| (config-vpls-sig)#exit | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v2 26 | Enter VPLS config mode |
| (config-vpls)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling ldp | Define Signaling as LDP |

| | |
|--|--------------------------------------|
| (config-vpls-sig)#vpls-type ethernet | Type ethernet configuration for VPLS |
| (config-vpls-sig)#vpls-peer 2.2.2.2 tunnel-name t2 | Configure VPLS Peer |
| (config-vpls-sig)#commit | Commit the transaction. |
| (config-vpls)#exit | Exit VPLS mode |

PE2: LDP VPLS Configuration

| | |
|--|---|
| (config)#mpls vpls v1 25 | Enter VPLS config mode |
| (config-vpls)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 1.1.1.1 tunnel-name t3 | Configure VPLS Peer |
| (config-vpls-sig)# exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v2 26 | Enter VPLS config mode |
| (config-vpls)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type ethernet | Type ethernet configuration for VPLS |
| (config-vpls-sig)#vpls-peer 1.1.1.1 tunnel-name t3 | Configure VPLS Peer with tunnel-name t2 |
| (config-vpls-sig)#commit | Commit the transaction. |
| (config-vpls)#exit | Exit VPLS mode |

LDP VPLS Service Mapping Configuration

PE1

| | |
|--|---|
| #configure terminal | Configure mode |
| (config)#service-template template1 | Template configuration |
| (config-svc)# match double-tag outer-vlan 2024 inner-vlan 2023 | Match criteria under template configuration |
| (config-svc)# rewrite ingress pop outgoing-tpid dot1.q | Action to be performed for the match. |
| (config-svc)#exit | Exit template configuration mode |
| (config)#service-template template4 | Template configuration |
| (config-svc)# match outer-vlan 700 | Allow VLAN 700 traffic on this VC |
| (config-svc)# match double-tag outer-vlan 1200 inner-vlan 3200 | Allow double tag match with s+c tags |
| (config-svc)# match untagged | Allow untagged traffic |
| (config-svc)# rewrite ingress push 300 | Push Action performed for service template |
| (config-svc)#commit | Commit the transaction. |

PE1: Access port Configuration

| | |
|---|---|
| (config)#interface xe15 | Enter the access interface xe15. |
| (config-if)#switchport | Configure interface as a layer 2 port. |
| (config-if)#mpls-vpls v1 service-template template1 | Bind the VPLS to the Access Interface. |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if-vpls)#exit-if-vpls | Exit VPLS attachment-circuit mode |
| (config-if)#mpls-vpls v2 service-template template4 | Bind the VPLS to the Access Interface. |
| (config-if-vpls)#commit | Commit the transaction. |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |

PE2

| | |
|---|---|
| #configure terminal | Configure mode |
| (config)#service-template template1 | Template configuration |
| (config-svc)# match double-tag outer-vlan 2024 inner-vlan 2023 | Match criteria under template configuration |
| (config-svc)# rewrite ingress pop outgoing-tpid dot1.q | Action to be performed for the match. |
| (config-svc)#exit | Exit template configuration mode |
| (config)#service-template template4 | Template configuration |
| (config-svc)# match outer-vlan 700 | Allow VLAN 700 traffic on this VC |
| (config-svc)# match double-tag outer-vlan 1200 inner-vlan 3200 | Allow double tag match with s+c tags |
| (config-svc)# match untagged | Allow untagged traffic |
| (config-svc)# rewrite ingress push 300 | Push Action performed for service template |
| (config-svc)#commit | Commit the transaction. |

PE2: Access port Configuration

| | |
|---|---|
| (config)#interface xe12 | Enter access Interface xe12 |
| (config-if)#switchport | Configure interface as a layer 2 port. |
| (config-if)#mpls-vpls v1 service-template template1 | Bind the VPLS to the Access Interface. |
| (config-if-vpls)#exit-if-vpls | Exit VPLS attachment-circuit mode |
| (config-if)#mpls-vpls v2 service-template template4 | Bind the VPLS to the Access Interface. |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if-vpls)#commit | Commit the transaction. |
| (config-if)#exit | Exit Interface mode and return to Configure mode. |

Validation

Below are the example outputs of mpls vpls with tunnel-name.


```

PE1#show mpls vpls mesh
VPLS-ID      Peer Addr      Tunnel-Label  In-Label      Network-Intf   Out-Label  Lkps/St  PW-INDEX
SIG-Protocol  Status
25           2.2.2.2           24321        24322         xe1            24322      2/Up     4          L
DP           Active
26           2.2.2.2           24321        24323         xe1            24323      2/Up     5          L
DP           Active
PE1#

PE1#show mpls vpls detail
Virtual Private LAN Service Instance: v1, ID: 25
SIG-Protocol: LDP
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 88a8
Configured interfaces:
  Interface: xe15
Service-template : template1
Match criteria : 2024/2023
Action type : Pop
Outgoing tpid : dot1.q

Mesh Peers:
  2.2.2.2 (Up)
  Tunnel-Name: t2

Virtual Private LAN Service Instance: v2, ID: 26
SIG-Protocol: LDP
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Raw
Configured interfaces:
  Interface: xe15
Service-template : template4
Match criteria : 700,
1200/3200,
Untagged
Action type : Push
Action value : 300

Mesh Peers:
  2.2.2.2 (Up)
  Tunnel-Name: t2

PE2#

PE1#sh mpls vpls mesh
VPLS-ID      Peer Addr      Tunnel-Label  In-Label      Network-Intf   Out-Label  Lkps/St  PW-INDEX
SIG-Protocol  Status
25           2.2.2.2           24321        24322         Xe1            24322      2/Up     4          L
DP           Active
26           2.2.2.2           24321        24323         Xe1            24323      2/Up     5          L
DP           Active

s
PE1#

```

```

PE2#sh mpls vpls detail
Virtual Private LAN Service Instance: v1, ID: 25 SIG-Protocol: LDP
Attachment-Circuit :UP Learning: Enabled Control-Word: Disabled
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500 Description: none
service-tpid: dot1ad Operating mode: Tagged Svlan Id: 0
Svlan Tpid: 88a8 Configured interfaces:
Interface: xe15
Service-template : template1 Match criteria : 2024/2023 Action type : Pop
Outgoing tpid : dot1q
Mesh Peers:
2.2.2.2 (Up) Tunnel-Name: t2

Virtual Private LAN Service Instance: v2, ID: 26 SIG-Protocol: LDP
Attachment-Circuit :UP Learning: Enabled Control-Word: Disabled
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500 Description: none
service-tpid: dot1ad Operating mode: Raw Configured interfaces:
Interface: xe15
Service-template : template4 Match criteria : 700,
1200/3200,
Untagged
Action type : Push Action value : 300
Mesh Peers:
2.2.2.2 (Up) Tunnel-Name: t2

```

Configure Static VPLS

PE1: LDP VPLS Configuration

| | |
|---|---|
| (config)#mpls vpls v3 27 | Enter VPLS config mode |
| (config-vpls)#vpls-peer 2.2.2.2 tunnel-name t2 manual | Configure VPLS Peer with trunk-name t2 with manual option |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#service-template vpls1 | Template configuration |
| (config-svc)# match outer-vlan 1000 | Allow VLAN 1000 traffic on this VC |
| (config-svc)#exit | Exit service template mode |

PE1: Access port Configuration

| | |
|---|--|
| (config)#interface xe15 | Enter the access Interface xe15 |
| (config-if)#switchport | Configure interface as a layer 2 port. |
| (config-if)#mpls-vpls v3 service-template vpls1 | Bind the VPLS to the Access Interface. |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if-vpls)#exit-if-vpls | Exit VPLS attachment-circuit mode |
| (config-if)#exit | Exit from the interface mode |
| (config)#vpls fib-entry 27 peer 2.2.2.2 3000 xe1 4000 | Configure VPLS FIB entry for VPLS peer PE2 |
| (config)#commit | Commit the transaction. |

PE2: LDP VPLS Configuration

| | |
|---|--|
| (config)#mpls vpls v3 27 | Enter VPLS config mode |
| (config-vpls)#vpls-peer 1.1.1.1 tunnel-name t3 manual | Configure static VPLS Peer with tunnel-name t3 |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#service-template vpls1 | Template configuration |
| (config-svc)# match outer-vlan 1000 | Allow VLAN 1000 traffic on this VC |
| (config-svc)#exit | Exit service template mode |

PE2: Access port Configuration

| | |
|--|---|
| (config)#interface xe12 | Enter the access interface xe12 |
| (config-if)#switchport | Configure interface as a layer 2 port. |
| (config-if)#mpls-vpls v3 service-template vpls1 | Bind the VPLS to the Access Interface. |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if-vpls)#exit-if-vpls | Exit VPLS attachment-circuit mode |
| (config-if)#exit | Exit interface mode. |
| (config)#vpls fib-entry 27 peer 1.1.1.1 4000 xe13 3000 | Configure VPLS FIB entry for VPLS peer PE1. |
| (config)#commit | Commit the transaction. |

Validation

```

PE1#show mpls vpls mesh
VPLS-ID    Peer Addr      Tunnel-Label  In-Label    Network-Intf  Out-Label    Lkps/St    PW-INDEX
SIG-Protocol Status
25         2.2.2.2         24321        24322       xe1           24322        2/Up       4          L
DP         Active
26         2.2.2.2         24321        24323       xe1           24323        2/Up       5          L
DP         Active
27         2.2.2.2         24321        3000        xe1           4000         2/Up       6          S
TATIC     Active

PE1#show mpls vpls v3 detail
Virtual Private LAN Service Instance: v3, ID: 27
SIG-Protocol: STATIC
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Configured interfaces:
  Interface: xe15
Service-template : vpls1
Match criteria : 1000

Mesh Peers:
  2.2.2.2 (Up)
  Tunnel-Name: t2

PE2#

```

```
PE2#sh mpls vpls mesh
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|---------|-----------|--------------|----------|--------------|-----------|---------|----------|
| 25 | 2.2.2.2 | 24321 | 24322 | xe1 | 24322 | 2/Up | 4 |
| DP | Active | | | | | | L |
| 26 | 2.2.2.2 | 24321 | 24323 | xe1 | 24323 | 2/Up | 5 |
| DP | Active | | | | | | L |
| 27 | 2.2.2.2 | 24321 | 3000 | xe1 | 4000 | 2/Up | 6 |
| TATIC | Active | | | | | | S |

```
PE2#
```

```
PE2#sh mpls vpls v3 detail
```

```
Virtual Private LAN Service Instance: v3, ID: 27 SIG-Protocol: STATIC
```

```
Attachment-Circuit :UP Learning: Enabled Control-Word: Disabled
```

```
Group ID: 0, Configured MTU: 1500 Description: none
```

```
service-tpid: dot1.q Operating mode: Raw Configured interfaces:
```

```
Interface: xe15
```

```
Service-template : vpls1 Match criteria : 1000
```

```
Mesh Peers:
```

```
2.2.2.2 (Up) Tunnel-Name: t2
```

MPLS BFD Configuration

This section explains how to configure MPLS-LSP BFD and PW-VCCV BFD.

Users can use BFD to monitor the liveness of MPLS LSPs. Additionally, BFD serves as a control and verification (CV) mechanism in VCCV to detect faults and signal status for pseudowires (PWs).

The Qumran (Q1 and Q2) series platforms support VCCV BFD with the following restrictions:

- Qumran (Q1 and Q2) series platforms: Support both CC-Type 1 and 2.
- Qumran2 (Q2) series platforms: Support only CC-Type 1 and CV-Type 1 and 3.




Note: Starting with OcNOS version 7.0.0, sub-interfaces replace Service Templates as the preferred configuration method. While legacy Service Template configurations remain functional, the associated commands are now hidden and will be completely removed in a future release. To ensure future-proofing and continued manageability, please migrate all existing Service Templates to sub-interfaces. For more details, see [Appendix A: Service Template Configuration \(page 2631\)](#).

Topology

The topology diagram illustrates the configuration needed to enable MPLS LSP BFD.

Figure 45. MPLS LSP BFD Topology



Configuring MPLS BFD LSP over LDP

PE1

1. **Configure Loopback Interface:** Set the loopback interface IP to 1.1.1.1/32.

```
# configure terminal
(config)# interface lo
(config-if)# ip address 1.1.1.1/32 secondary
```

2. **Configure Ethernet Interface**

- Assign IP address to eth1.

- Enable label switching and LDP on the interface.

```
(config)# interface eth1
(config-if)# ip address 10.10.10.1/24
(config-if)# label-switching
(config-if)# enable-ldp ipv4
(config-if)# exit
```

3. Configure LDP: Set the LDP router ID and transport address (preferably the loopback address).

```
(config)# router ldp
(config-router)# router-id 1.1.1.1
(config-router)# transport-address ipv4 1.1.1.1
(config-router)# exit
```

4. Configure OSPF

- Start the OSPF process (process ID 100).
- Advertise the eth1 and loopback interfaces in Area 0.

```
(config)# router ospf 100
(config-router)# network 10.10.10.0/24 area 0
(config-router)# network 1.1.1.1/32 area 0
(config-router)# exit
```

5. Configure MPLS BFD for LDP Peers: Enable MPLS BFD for LDP peers 2.2.2.2 and 3.3.3.3. Save and apply the configuration.

```
(config)# mpls bfd ldp 2.2.2.2/32
(config-mpls-bfd)# exit

(config)# mpls bfd ldp 3.3.3.3/32
(config-mpls-bfd)# exit

(config)# commit
(config)# end
```

P

1. Configure Loopback Interface: Set the loopback interface IP to 2.2.2.2/32.

```
# configure terminal
(config)# interface lo
(config-if)# ip address 2.2.2.2/32 secondary
```

2. Configure Ethernet Interfaces

- Assign IP address to eth1 and eth2 interfaces.
- Enable label switching and LDP on the interfaces eth1 and eth2.

```
(config)# interface eth1
(config-if)# ip address 10.10.10.2/24
(config-if)# label-switching
(config-if)# enable-ldp ipv4
(config-if)# exit

(config)# interface eth2
(config-if)# ip address 20.20.20.1/24
(config-if)# label-switching
(config-if)# enable-ldp ipv4
(config-if)# exit
```

3. Configure LDP: Set the LDP router ID and transport address (preferably the loopback address).

```
(config)# router ldp
(config-router)# router-id 2.2.2.2
(config-router)# transport-address ipv4 2.2.2.2
(config-router)# exit
```

4. Configure OSPF

- Start the OSPF process (process ID 100).
- Advertise the eth1, eth2, and loopback interfaces in Area 0.

```
(config)# router ospf 100
(config-router)# network 10.10.10.0/24 area 0
(config-router)# network 20.20.20.0/24 area 0
(config-router)# network 2.2.2.2/32 area 0
(config-router)# exit
```

5. Configure MPLS BFD for LDP Peers: Enable MPLS BFD for LDP peers 1.1.1.1 and 3.3.3.3. Save and apply the configuration.

```
(config)# mpls bfd ldp 3.3.3.3/32
(config-mpls-bfd)# exit

(config)# mpls bfd ldp 1.1.1.1/32
(config-mpls-bfd)# exit

(config)# commit
(config)# end
```

PE2**1. Configure Loopback Interface:** Set the loopback interface IP to 3.3.3.3/32.

```
# configure terminal
(config)# interface lo
(config-if)# ip address 3.3.3.3/32 secondary
```

2. Configure Ethernet Interface

- Assign IP address to eth2.
- Enable label switching and LDP on the interface.

```
(config)# interface eth2
(config-if)# ip address 20.20.20.2/24
(config-if)# label-switching
(config-if)# enable-ldp ipv4
(config-if)# commit
(config-if)# exit
```

3. Configure LDP: Set the LDP router ID and transport address (preferably the loopback address).

```
(config)# router ldp
(config-router)# router-id 3.3.3.3
(config-router)# transport-address ipv4 3.3.3.3
(config-router)# exit
```

4. Configure OSPF

- Start the OSPF process (process ID 100).

- Advertise the eth2 and loopback interfaces in Area 0.

```
(config)# router ospf 100
(config-router)# network 20.20.20.0/24 area 0
(config-router)# network 3.3.3.3/32 area 0
(config-router)# commit
(config-router)# exit
```

5. Configure MPLS BFD for LDP Peers: Enable MPLS BFD for LDP peers 2.2.2.2 and 1.1.1.1. Save and apply the configuration.

```
(config)# mpls bfd ldp 2.2.2.2/32
(config-mpls-bfd)# exit

(config)# mpls bfd ldp 1.1.1.1/32
(config-mpls-bfd)# exit

(config)# commit
(config)# end
```

Validation

OSPF Neighbor Adjacency Verification

Use the following command outputs to verify that OSPF neighbor adjacency has been successfully established between all five routers.

```
PE1#show ip ospf neighbor
```

```
Total number of full neighbors: 1
```

```
OSPF process 1 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|------------|-----------|-------------|
| 2.2.2.2 | 1 | Full/DR | 00:00:32 | 10.10.10.2 | eth1 | 0 |

```
P#show ip ospf neighbor
```

```
Total number of full neighbors: 2
```

```
OSPF process 1 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 1.1.1.1 | 1 | Full/Backup | 00:00:35 | 10.10.10.1 | eth1 | 0 |
| 3.3.3.3 | 1 | Full/Backup | 00:00:32 | 20.20.20.2 | eth2 | 0 |

```
PE2#show ip ospf neighbor
```

```
Total number of full neighbors: 1
```

```
OSPF process 1 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|------------|-----------|-------------|
| 2.2.2.2 | 1 | Full/DR | 00:00:31 | 20.20.20.1 | eth1 | 0 |

LDP Session Verification

The outputs below confirm that LDP sessions are operational on all three routers involved in the MPLS topology.

```
PE1#show ldp session
```

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 2.2.2.2 | eth1 | Passive | OPERATIONAL | 30 | 01:30:57 |

```
P#show ldp session
```

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 1.1.1.1 | eth1 | Active | OPERATIONAL | 30 | 01:48:34 |
| 3.3.3.3 | eth2 | Passive | OPERATIONAL | 30 | 03:01:14 |

```
PE2#show ldp session
```


| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 2.2.2.2 | eth2 | Active | OPERATIONAL | 30 | 03:01:14 |

BFD MPLS LSP Session Verification

The command results below show that BFD sessions over MPLS LSPs using LDP as the signaling protocol are up and functioning correctly.

```
PE1# show bfd mpls lsp sessions
Peer Addr      Path      Tunnel-
name          State    Local    Local  Min      Min      Encap    Last      Multi  T
                                Disc
x             Rx             Dn-Time
2.2.2.2       LDP
LSP           Up      1281     3       10000    10000    IP/UDP   00:00
:00
3.3.3.3       LDP
LSP           Up      1282     3       10000    10000    IP/UDP   00:00
:00

P #show bfd mpls lsp sessions
Peer Addr      Path      Tunnel-
name          State    Local    Local  Min      Min      Encap    Last      Multi  T
                                Disc
x             Rx             Dn-Time
3.3.3.3       LDP
LSP           Up      1283     3       10000    10000    IP/UDP   00:00
:00
1.1.1.1       LDP
LSP           Up      1282     3       10000    10000    IP/UDP   00:00
:00

PE1 #show bfd mpls lsp sessions
Peer Addr      Path      Tunnel-
name          State    Local    Local  Min      Min      Encap    Last      Multi  T
                                Disc
x             Rx             Dn-Time
2.2.2.2       LDP
LSP           Up      1282     3       10000    10000    IP/UDP   00:00
:00
1.1.1.1       LDP
LSP           Up      1281     3       10000    10000    IP/UDP   00:00
:00
```

Configuring MPLS BFD LSP over RSVP

This section explains the configuration of an RSVP-based MPLS BFD session between PE1 and PE2.



Note: For OSPF configurations, refer to the [Configuring MPLS BFD LSP over LDP \(page 1289\)](#) section.

PE1

1. Configure RSVP and Ethernet Interface

- Configure RSVP router
- Assign IP address to eth1.
- Enable label switching and RSVP on the interface eth1.

```
(config)#router rsvp
```

```
(config-router)#exit
(config)#interface eth1
(config-if)#label-switching
(config-if)#enable-rsvp
(config-if)#exit
```

2. Configure RSVP Tunnel: Configure the RSVP tunnel (T1) from PE1 to PE2.

```
(config)# rsvp-trunk T1 ipv4
(config-trunk)# from 1.1.1.1
(config-trunk)# to 3.3.3.3
(config-trunk)# exit
```

3. Configure MPLS BFD for RSVP Tunnel: Enable MPLS BFD for RSVP tunnel T1. Save and apply the configuration.

```
(config)# mpls bfd rsvp tunnel-name T1
(config-mpls-bfd)# exit

(config)# commit
(config)# end
```

P

Configure RSVP and Ethernet Interfaces

- Configure RSVP router
- Assign IP address to `eth1` and `eth2` interfaces.
- Enable label switching and RSVP on the interfaces `eth1` and `eth2`.

```
(config)#router rsvp
(config-router)#exit

(config)#interface eth1
(config-if)#label-switching
(config-if)#enable-rsvp
(config-if)#exit

(config)# interface eth2
(config-if)# label-switching
(config-if)# enable-rsvp
(config-if)# exit

(config)# commit
(config)# exit
```

PE2

1. Configure RSVP and Ethernet Interface

- Configure RSVP router
- Assign IP address to `eth2`.
- Enable label switching and RSVP on the interface `eth2`.

```
(config)#router rsvp
(config-router)#exit
(config)#interface eth2
(config-if)#label-switching
(config-if)#enable-rsvp
(config-if)#exit
```

2. Configure RSVP Tunnel: Configure the RSVP tunnel (T1) from PE2 to PE1.

```

PE2(config)# rsvp-trunk T1 ipv4
PE2(config-trunk)# from 3.3.3.3
PE2(config-trunk)# to 1.1.1.1
PE2(config-trunk)# exit

```

3. Configure MPLS BFD for RSVP Tunnel: Enable MPLS BFD for RSVP tunnel T1. Save and apply the configuration.

```

(config)# mpls bfd rsvp tunnel-name T1
(config-mpls-bfd)# exit

(config)# commit
(config)# end

```

Validation

Verify that the RSVP signaling session between PE1 and PE2 is established and operational.

```

PE1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

Ingress RSVP:

| To | From | Type | LSPName | State |
|---------|---------|---------|------------------|--------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 3.3.3.3 | 1.1.1.1 | PRI | T1-Primary | UP |
| | 24321 | DEFAULT | | |

Total 1 displayed, Up 1, Down 0.

Egress RSVP:

| To | From | Type | LSPName | State |
|---------|---------|-------|------------------|--------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 1.1.1.1 | 3.3.3.3 | PRI | T1-Primary | UP |
| SE | 24320 | - | ELSP_CON | |

Total 1 displayed, Up 1, Down 0.

PE2#show rsvp session

```

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

Ingress RSVP:

| To | From | Type | LSPName | State |
|---------|---------|---------|------------------|--------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 1.1.1.1 | 3.3.3.3 | PRI | T1-Primary | UP |
| | 24320 | DEFAULT | | |

Total 1 displayed, Up 1, Down 0.

Egress RSVP:

| To | From | Type | LSPName | State |
|---------|---------|-------|------------------|--------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 3.3.3.3 | 1.1.1.1 | PRI | T1-Primary | UP |
| SE | 24320 | - | ELSP_CON | |

Confirm that MPLS BFD sessions over the RSVP tunnel between PE1 and PE2 are up and functioning correctly.

PE1#show bfd mpls lsp sessions

| Peer Addr | Path | Tunnel- | State | Local | Local | Min | Min | Encap | Last | Multi | T |
|-----------|------|---------|-------|-------|-------|-------|-------|--------|--------|-------|---|
| name | | | | | | | | | Disc | | |
| x | Rx | Dn-Time | | | | | | | | | |
| 3.3.3.3 | | RSVP | | | | | | | | | |
| LSP T1 | | | Up | 1282 | 3 | 10000 | 10000 | IP/UDP | 00:00: | | |

```

00
PE2#show bfd mpls lsp sessions
Peer Addr          Path          Tunnel-
name              State    Local    Local    Min        Min        Encap    Last    Multi  T
                                Disc
x                Rx                Dn-Time
1.1.1.1          RSVP
LSP   T1                Up        1281      3        10000      10000      IP/UDP  00:00:
00

```

Configuring Static MPLS BFD

This section shows how to configure MPLS BFD statically.



Note: For OSPF configurations, please refer 'Configuring MPLS BFD LSP over LDP' section.

PE1

| | |
|--|---|
| (config)#interface eth1 | Enter interface mode |
| (config-if)#label-switching | Enable label switching on interface eth1. |
| (config-if)#exit | Exit interface mode |
| (conf)#mpls ftn-entry 3.3.3.3/32 102 10.10.10.2 eth1 | Configure MPLS FTN entry for the creation of a static LSP to |
| PE-2. | |
| (config)#mpls ilm-entry 401 pop | Configure ILM entry |
| (config)#mpls bfd static all | Configure mpls bfd for static MPLS |
| (config-mpls-bfd)#exit | Exit mpls bfd mode |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | Exit out of configuration terminal mode |

P

| | |
|--|---|
| (config)#interface eth1 | Enter interface mode |
| (config-if)#label-switching | Enable label switching on interface eth1. |
| (confi-if)#exit | Exit interface mode |
| (conf)#interface eth2 | Enter interface mode |
| (config-if)#label-switching | Enable label switching on interface eth2. |
| (config-if)#exit | Exit interface mode |
| mpls ilm-entry 201 swap 401 eth1 10.10.10.1 1.1.1.1/32 | Swap the incoming label |
| mpls ilm-entry 102 swap 301 eth2 mpls ftn-entry 1.1.1.1/32 201 20.20.20.1 eth2 20.20.20.2 3.3.3.3/32 | Swap the incoming label |

| | |
|------------------------------|---|
| (config)#mpls bfd static all | Configure mpls bfd for static MPLS |
| (config-mpls-bfd)#exit | Exit mpls bfd mode |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | Exit out of configuration terminal mode |

PE2

| | |
|--|---|
| (conf)#interface eth2 | Enter interface mode |
| (config-if)#label-switching | Enable label switching on interface eth2. |
| (config-if)#exit | Exit interface mode |
| (config)#mpls ftn-entry 1.1.1.1/32 201 20.20.20.1 eth2 | Configure MPLS FTN entry for the creation of a static LSP to |
| PE-1 | |
| (config)#mpls bfd static all | Configure mpls bfd for static MPLS |
| (config-mpls-bfd)#exit | Exit mpls bfd mode |
| (config)#commit | Commit the candidate configuration to the running configuration |
| (config)#end | Exit out of configuration terminal mode |

Validations

Verify static MPLS BFD session is up and running between PE nodes.

```

PE1#show bfd mpls lsp sessions
Peer Addr      Path      Tunnel-
name           State    Local    Local    Min      Min      Encap    Last    Multi  T
              Disc
x             Rx      Dn-Time
3.3.3.3        Static
LSP            Up      1281     3        10000    10000    IP/UDP   00:00:00

PE2#show bfd mpls lsp sessions
Peer Addr      Path      Tunnel-
name           State    Local    Local    Min      Min      Encap    Last    Multi  T
              Disc
x             Rx      Dn-Time
1.1.1.1        Static
LSP            Up      1281     3        10000    10000    IP/UDP   00:00:00

```

Configuring PW VCCV BFD

This section shows how to configure PW VCCV BFD.



Note: For OSPF configurations, refer to [Configuring MPLS BFD LSP over LDP \(page 1289\)](#) section.

PE1

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| (conf)#interface eth1 | Enter interface mode |
| (conf-if)#ip address 10.10.10.1/24 | Assign IP address to eth1 interface |
| (config-if)#label-switching | Enable label switching on interface eth1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on eth1. |
| (confi-if)#exit | Exit interface mode |
| (config)#router ldp | Enter Router mode for LDP. |
| (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1 |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter "ipv6" if you are configuring an IPv6 interface. |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#service-template st1 | Service template configuration |
| (config-svc)#exit | Exit service-template configuration mode |
| (config)#mpls l2-circuit pw1 3.3.3.3 | Configure the VC for PE2 |
| (config-pseudowire)#control-word | Configure control-word |
| (config-pseudowire)#vccv cc-type type-1 | Configure cc-type type-1 |
| (config-pseudowire)#vccv cv-type type-3 | Configure cv-type type-3 |
| (confi- pseudowire)#exit | Exit pw configuration mode |
| (config)#commit | Commit the candidate configuration to the running configuration |

P

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface eth1 | Specify the interface (eth1) to be configured. |
| (conf-if)#ip address 10.10.10.2/24 | Assign IP address to eth1 interface |
| (config-if)#label-switching | Enable label switching on interface eth1. |

| | |
|--|--|
| (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface (eth1). |
| (config-if)#exit | Exit interface mode. |
| (config)#interface eth2 | Specify the interface (eth2) to be configured. |
| (conf-if)#ip address 20.20.20.1/24 | Assign IP address to eth2 interface |
| (config-if)#label-switching | Enable label switching on interface eth2. |
| (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface (eth2). |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter Router mode. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2. |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#pw-status-tlv | Configure PW Status TLV to signal the pseudowire status |
| (config-router)#exit | Exit Router mode and return to Configure mode. |
| (config)#commit | Commit the candidate configuration to the running configuration |

PE2

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 3.3.3.3/32 secondary | Set the IP address of the loopback interface to 3.3.3.3/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface eth1 | Specify the interface (eth1) to be configured. |
| (conf-if)#ip address 20.20.20.2/24 | Assign IP address to eth1 interface |
| (config-if)#label-switching | Enable label switching on interface eth2. |
| (config-if)#enable-ldp ipv4 | Enable LDP on a specified interface (eth2). |
| (config-if)#commit | Commit the transaction. |
| (config-if)#exit | Exit Interface mode |
| (config)#router ldp | Enter Router mode. |
| (config-router)#router-id 3.3.3.3 | Set the router ID for IP address 3.3.3.3. |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter "ipv6" if you are configuring an IPv6 interface. |
| (config-router)#exit | Exit the Router mode and return to the Configure mode. |
| (config)#service-template st1 | Service template configuration |

| | |
|---|--|
| (config-svc)#exit | Exit service-template configuration mode |
| (config)#mpls l2-circuit pw1 1.1.1.1 | Configure the VC for PE1 |
| (config-pseudowire)#control-word | Configure control-word |
| (config-pseudowire)#vccv cc-type type-1 | Configure cc-type type-1 |
| (config-pseudowire)#vccv cv-type type-3 | Configure cv-type type-3 |
| (confi- pseudowire)#exit | Exit pw configuration mode |

Validation

Verify PW VCCV BFD sessions are up between PE nodes.

```
PE1#show bfd mpls pw-vccv sessions
```

| VC-ID | Peer | | State | Local | Local | Min | Min | | Encap | Last | |
|-------|----------|--|-------|-------|-------|-------|------|--|-------|------|-----|
| Addr | | | | | Disc | Multi | Tx | | Rx | | Dn- |
| Time | | | | | | | | | | | |
| 1 | 3.3.3.3 | | Up | | 1281 | 3 | 3000 | | 3000 | PW- | |
| ACH | 00:00:00 | | | | | | | | | | |

```
PE2#show bfd mpls pw-vccv sessions
```

| VC-ID | Peer | | State | Local | Local | Min | Min | | Encap | Last | |
|-------|----------|--|-------|-------|-------|-------|------|--|-------|------|-----|
| Addr | | | | | Disc | Multi | Tx | | Rx | | Dn- |
| Time | | | | | | | | | | | |
| 1 | 1.1.1.1 | | Up | | 1281 | 3 | 3000 | | 3000 | PW- | |
| ACH | 00:00:00 | | | | | | | | | | |


BGP Peer Groups for Address-Family L2VPN EVPN

BGP peer groups are used to simplify configuration and to improve performance. This is achieved by assigning the same outbound policy to each of the neighbors. Because UPDATES are generated only once per peer group rather than multiple times for each neighboring router, peer groups save processing time when building neighbor updates. It reduces the amount of system resources (CPU and memory) necessary in an update generation.

A BGP peer group reduces the load on system resources by allowing the routing table to be checked only once, and updates to be replicated to all peer group members instead of being done individually for each peer in the peer group.

Topology

Figure 46. BGP Peer-Groups with L2VPN EVPN address-family



Configuration

R1

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)# interface lo | Enter interface mode for Loopback |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure ip address for Loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)# interface xe15 | Enter interface mode for xe15 |
| (config-if)#ip address 10.1.1.2/24 | Configure ip address |
| (config-if)#exit | Exit interface mode |

| | |
|--|---|
| (config)# interface ce0 | Enter interface mode for ce0 |
| (config-if)#ip address 20.1.1.2 | Configure ip address |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf 100 | Configure the OSPF process (100) |
| (config-router)# ospf router-id 1.1.1.1 | Configure OSPF router-id |
| (config-router)#network 1.1.1.1/32 area 0 | Advertise the network in Area 0 |
| (config-router)#network 10.1.1.0/24 area 0 | Advertise the network in Area 0 |
| (config-router)#network 20.1.1.0/24 area 0 | Advertise the network in Area 0 |
| (config-router)#exit | Exit Router mode and return to Configure mode |
| (config)#router bgp 100 | Define the routing process. The number 100 specifies the AS number of R1. |
| (config-router)# bgp router-id 1.1.1.1 | Configure BGP router-id |
| (config-router)#neighbor PG peer-group | Create a peer group named PG |
| (config-router)#neighbor PG remote-as 100 | Assign options to the peer group named PG |
| (config-router)#neighbor PG update-source lo | Assign options to the peer group named PG |
| (config-router)#neighbor 2.2.2.2 peer-group PG | Define neighbor 2.2.2.2 (R2) as a peer group |
| (config-router)#neighbor 3.3.3.3 peer-group PG | Define neighbor 3.3.3.3 (R3) as a peer group member. |
| (config-router)#address-family l2vpn evpn | Enter address-family l2vpn evpn mode |
| (config-router-af)#neighbor PG activate | Activate the peer-group ABC for address-family l2vpn evpn |
| (config-router-af)# exit-address-family | Exit address-family ipv4 unicast mode |
| (config-router)#exit | Exit router bgp mode |
| (config)#commit | Commit the candidate configuration to the running configuration. |

R2

| | |
|------------------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)# interface lo | Enter interface mode for Loopback |
| | Configure ip address for Loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)# interface xe15 | Enter interface mode for xe15 |
| (config-if)#ip address 10.1.1.2/24 | Configure ip address |
| (config-if)#exit | Exit interface mode |
| (config)# interface xe10 | Enter interface mode for xe10 |
| | Configure ip address |
| (config-if)#exit | Exit interface mode |

| | |
|--|---|
| (config)#router ospf 100 | Configure the OSPF process (100) |
| (config-router)# ospf router-id 2.2.2.2 | Configure OSPF router-id |
| (config-router)#network 2.2.2.2/32 area 0 | Advertise the network in Area 0 |
| (config-router)#network 10.1.1.0/24 area 0 | Advertise the network in Area 0 |
| (config-router)#exit | Exit Router mode and return to Configure mode |
| (config)#router bgp 100 | Define the routing process. The number 100 specifies the AS number of R1. |
| (config-router)# bgp router-id 2.2.2.2 | Configure BGP router-id |
| (config-router)#neighbor PG peer-group | Create a peer group named PG |
| (config-router)#neighbor PG remote-as 100 | Assign options to the peer group named PG |
| (config-router)#neighbor PG update-source lo | Assign options to the peer group named PG |
| (config-router)#neighbor 1.1.1.1 peer-group PG | Define neighbor 1.1.1.1 (R1) as a peer group member. |
| (config-router)#address-family l2vpn evpn | Enter address-family l2vpn evpn mode |
| (config-router-af)#neighbor PG activate | Activate the peer-group ABC for address-family l2vpn evpn |
| (config-router-af)# exit-address-family | Exit address-family ipv4 unicast mode |
| (config-router)#exit | Exit router bgp mode |
| (config)#commit | Commit the candidate configuration to the running configuration. |

R3

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)# interface lo | Enter interface mode for Loopback |
| | Configure ip address for Loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)# interface ce15 | Enter interface mode for ce15 |
| | Configure ip address |
| (config-if)#exit | Exit interface mode |
| (config)# interface xe10 | Enter interface mode for xe10 |
| | Configure ip address |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf 100 | Configure the OSPF process (100) |
| (config-router)# ospf router-id 3.3.3.3 | Configure OSPF router-id |
| (config-router)#network 20.1.1.0/24 area 0 | Advertise the network in Area 0 |
| (config-router)#exit | Exit Router mode and return to Configure mode |
| (config)#router bgp 100 | Define the routing process. The number 100 specifies the AS number of R1. |

| | |
|--|--|
| (config-router)# bgp router-id 3.3.3.3 | Configure BGP router-id |
| (config-router)#neighbor PG peer-group | Create a peer group named PG |
| (config-router)#neighbor PG remote-as 100 | Assign options to the peer group named PG |
| (config-router)#neighbor PG update-source lo | Assign options to the peer group named PG |
| (config-router)#neighbor 1.1.1.1 peer-group PG | Define neighbor 1.1.1.1 (R1) as a peer group member. |
| (config-router)#address-family l2vpn evpn | Enter address-family l2vpn evpn mode |
| (config-router-af)#neighbor PG activate | Activate the peer-group ABC for address-family l2vpn evpn |
| (config-router-af)# exit-address-family | Exit address-family ipv4 unicast mode |
| (config-router)#exit | Exit router bgp mode |
| (config)#commit | Commit the candidate configuration to the running configuration. |

Validation

R1

```

R1#sh run bgp
!
router bgp 100
  bgp router-id 1.1.1.1
  neighbor PG peer-group
  neighbor PG remote-as 100
  neighbor PG update-source lo
  neighbor 2.2.2.2 peer-group PG
  neighbor 3.3.3.3 peer-group PG
!
  address-family l2vpn evpn
  neighbor PG activate
  exit-address-family
R1#sh bgp neighbors
BGP neighbor is 2.2.2.2, remote AS 100, local AS 100, internal link
Member of peer-group PG for session parameters
  BGP version 4, local router ID 1.1.1.1, remote router ID 2.2.2.2
  BGP state = Established, up for 01:20:53
  Last read 00:00:24, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family L2VPN EVPN: advertised and received
  Received 192 messages, 0 notifications, 0 in queue
  Sent 191 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
For address family: L2VPN EVPN
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  PG peer-group member
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  Accepted AD:0 MACIP:0 MCAST:0 ESI:0 PREFIX:0
  0 announced prefixes

Connections established 1; dropped 0

```

```

Local host: 1.1.1.1, Local port: 42981
Foreign host: 2.2.2.2, Foreign port: 179
Nexthop: 1.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 3.3.3.3, remote AS 100, local AS 100, internal link
Member of peer-group PG for session parameters
  BGP version 4, local router ID 1.1.1.1, remote router ID 3.3.3.3
  BGP state = Established, up for 01:36:13
  Last read 00:00:08, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family L2VPN EVPN: advertised and received
  Received 227 messages, 0 notifications, 0 in queue
  Sent 229 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
For address family: L2VPN EVPN
  BGP table version 1, neighbor version 1
  Index 3, Offset 0, Mask 0x8
  PG peer-group member
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  Accepted AD:0 MACIP:0 MCAST:0 ESI:0 PREFIX:0
  0 announced prefixes

Connections established 1; dropped 0
Local host: 1.1.1.1, Local port: 179
Foreign host: 3.3.3.3, Foreign port: 32857
Nexthop: 1.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
R1#sh ip ospf neighbor

```

```

Total number of full neighbors: 2
OSPF process 100 VRF(default):
Neighbor ID      Pri   State           Dead Time   Address        Interface       Instance ID
2.2.2.2          1    Full/Backup     00:00:38   10.1.1.1       xe15            0
3.3.3.3          1    Full/Backup     00:00:34   20.1.1.1       ce0            0
R1#sh bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | TblVer | InQ | OutQ | V | AS | MsgRcv | MsgSen | AD | MACIP | MCAST | ESI | PREFIX-ROUTE | Instance ID |
|----------|--------|-----|------|---|-----|--------|--------|----|-------|-------|-----|--------------|-------------|
| 2.2.2.2 | | | | 4 | 100 | 193 | 191 | | 1 | 0 | 0 | 01:21:07 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | | | | | | | | |
| 3.3.3.3 | | | | 4 | 100 | 227 | 229 | | 1 | 0 | 0 | 01:36:27 | 0 |
| 0 | 0 | 0 | 0 | 0 | 0 | | | | | | | | |

```
Total number of neighbors 2
```

```
Total number of Established sessions 2
```

R2

```

R2#sh run bgp
!
router bgp 100
  bgp router-id 2.2.2.2

```

```

neighbor PG peer-group
neighbor PG remote-as 100
neighbor PG update-source lo
neighbor 1.1.1.1 peer-group PG
!
address-family l2vpn evpn
neighbor PG activate
exit-address-family
!
R2#sh bgp neighbors
BGP neighbor is 1.1.1.1, remote AS 100, local AS 100, internal link
Member of peer-group PG for session parameters
  BGP version 4, local router ID 2.2.2.2, remote router ID 1.1.1.1
  BGP state = Established, up for 01:20:42
  Last read 00:00:20, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family L2VPN EVPN: advertised and received
  Received 190 messages, 0 notifications, 0 in queue
  Sent 193 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
For address family: L2VPN EVPN
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  PG peer-group member
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  Accepted AD:0 MACIP:0 MCAST:0 ESI:0 PREFIX:0
  0 announced prefixes

Connections established 1; dropped 0
Local host: 2.2.2.2, Local port: 179
Foreign host: 1.1.1.1, Foreign port: 42981
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

```

```
R2#sh ip ospf neighbor
```

```

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface       Instance ID
1.1.1.1        1     Full/DR         00:00:30   10.1.1.2       xe10            0

```

```

R2#sh bgp l2vpn evpn summary
BGP router identifier 2.2.2.2, local AS number 100
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | TblVer | InQ | OutQ | V | AS | MsgRcv | MsgSen | AD | MACIP | MCAST | ESI | PREFIX-ROUTE | Instance ID |
|----------|--------|-----|------|---------|----|--------|--------|-----|-------|-------|-----|--------------|-------------|
| 1.1.1.1 | 0 | 0 | 0 | Up/Down | 4 | 100 | 192 | 195 | 1 | 0 | 0 | 01:21:28 | 0 |

```
Total number of neighbors 1
```

```
Total number of Established sessions 1
```

R3

```

R3#sh run bgp
!
router bgp 100

```

```

bgp router-id 3.3.3.3
neighbor PG peer-group
neighbor PG remote-as 100
neighbor PG update-source lo
neighbor 1.1.1.1 peer-group PG
!
address-family l2vpn evpn
neighbor PG activate
exit-address-family
!
R3#sh bgp neighbors
BGP neighbor is 1.1.1.1, remote AS 100, local AS 100, internal link
Member of peer-group PG for session parameters
  BGP version 4, local router ID 3.3.3.3, remote router ID 1.1.1.1
  BGP state = Established, up for 01:36:07
  Last read 00:00:06, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family L2VPN EVPN: advertised and received
  Received 228 messages, 0 notifications, 0 in queue
  Sent 227 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
For address family: L2VPN EVPN
  BGP table version 1, neighbor version 1
  Index 2, Offset 0, Mask 0x4
  PG peer-group member
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  Accepted AD:0 MACIP:0 MCAST:0 ESI:0 PREFIX:0
  0 announced prefixes

Connections established 1; dropped 0
Local host: 3.3.3.3, Local port: 32857
Foreign host: 1.1.1.1, Foreign port: 179
Nexthop: 3.3.3.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
R3#sh ip os neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface       Instance ID
1.1.1.1        1     Full/DR         00:00:37    20.1.1.2       ce15            0
R3#sh bgp l2vpn evpn summary
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor      V   AS   MsgRcv   MsgSen   TblVer   InQ   OutQ   Up/Down   State/PfxRcd   AD   MACI
P MCAST      ESI  PREFIX-ROUTE
1.1.1.1      4   100    232      231      1        0      0      01:37:55      0      0
0            0    0        0        0        0        0      0
Total number of neighbors 1

Total number of Established sessions 1

```


Point-to-Point Connections Over MPLS

Overview

An MPLS Layer 2 Virtual Circuit (VC) facilitates efficient point-to-point Layer 2 connectivity in a service providers MPLS network. It enables Layer 2 circuit transport across the providers infrastructure, ensuring secure communication between remote sites through a single Label Switched Path (LSP) tunnel connecting Provider Edge (PE) routers. This feature optimizes network performance and supports diverse applications and services.

Topology

Figure 47. Point-to-Point Connection Over MPLS



The VC configuration process can be divided into the following steps:



Note: Loopback addresses being used should be advertised through OSPF, or must be statically routed.

1. Configure the IP address and OSPF for the PE1, P (Provider), and PE2 routers.
2. Configure MPLS and LDP on PE1, P, and PE2, and LDP targeted peer for the PE1 and PE2 routers. (If RSVP is used for configuring trunks, LDP must be configured on PE1 and PE2, and RSVP must be configured on PE1, P, and PE2.)
3. Configure the VC.
4. Bind the customer interface to the VC.

Configure IP Address and OSPF on Routers

Configure the IP addresses and OSPF on the PE1, P1, P2, PE2 and PE3 routers.

PE1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/ 32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe6 | Specify the interface (xe6) to be configured. |
| (config-if)#ip address 40.1.1.1/24 | Set the IP address of the interface to 40.1.1.1/24 |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 20.20.20.1/24 | Set the IP address of the interface to 20.20.20.1/24 |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 40.1.1.0/24 area 0 (config-router)#network 1.1.1.1/32 area 0 (config-router)# network 20.20.20.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |

P1

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 4.4.4.4/32 secondary | Set the IP address of the loopback interface to 9.9.9.9/ 32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe1/1 | Specify the interface (xe1/1) to be configured. |
| (config-if)#ip address 40.1.1.2/24 | Set the IP address of the interface to 40.1.1.2/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe33 | Specify the interface (xe33) to be configured. |
| (config-if)#ip address 50.1.1.2/24 | Set the IP address of the interface to 50.1.1.1/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 40.1.1.0/24 area 0 (config-router)#network 50.1.1.0/24 area 0 (config-router)#network 4.4.4.4/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |

P2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 5.5.5.5/32 secondary | Set the IP address of the loopback interface to 5.5.5.5/ 32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe6 | Specify the interface (xe1/1) to be configured. |
| (config-if)#ip address 70.1.1.1/24 | Set the IP address of the interface to 70.1.1.1/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe47 | Specify the interface (xe47) to be configured. |
| (config-if)#ip address 20.20.20.2/24 | Set the IP address of the interface to 20.20.20.2/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 20.20.20.20/24 area 0 (config-router)#network 70.1.1.0/24 area 0 (config-router)#network 5.5.5.5/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |

PE2

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/ 32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe33 | Specify the interface (xe33) to be configured. |
| (config-if)#ip address 50.1.1.1/24 | Set the IP address of the interface to 50.1.1.1/24 |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 50.1.1.0/24 area 0 (config-router)#network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs, and associate the area ID (0) with the interface. |

PE3

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 3.3.3.3/32 secondary | Set the IP address of the loopback interface to 3.3.3.3/ 32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe6 | Specify the interface (xe33) to be configured. |
| (config-if)#ip address 70.1.1.2/24 | Set the IP address of the interface 70.1.1.2/24 |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 70.1.1.0/24 area 0 (config-router)#network 3.3.3.3/32 area 0 | Define the interface on which OSPF runs, and associate the area ID (0) with the interface. |

Configure MPLS, LDP, and LDP Targeted Peer on Routers

PE1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter the Router mode. |
| (config)# router-id 1.1.1.1 | Configure LDP router ID. |
| (config-router)#pw-status-tlv | Set PW status TLV |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Specify the targeted LDP peer on PE1. |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Specify the targeted LDP peer on PE1. |
| (config-router-targeted-peer)# exit | Exit the Router targeted peer mode. |
| (config-router)#exit | Exit the Router mode. |
| (config)#interface xe1 | Specify the interface (xe1) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe1. |
| (config)#interface xe6 | Specify the interface (xe6) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe6 |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe6. |

PE2

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter the Router mode. |
| (config)# router-id 2.2.2.2 | Configure LDP router ID. |
| (config-router)#pw-status-tlv | Set PW status TLV |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Specify the targeted LDP peer on PE1. |
| (config-router-targeted-peer)# exit | Exit the Router targeted peer mode. |
| (config-router)#exit | Exit the Router mode. |
| (config)#interface xe33 | Specify the interface (xe33) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe33. |

PE3

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter the Router mode. |
| (config)# router-id 3.3.3.3 | Configure LDP router ID. |
| (config-router)#pw-status-tlv | Set PW status TLV |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Specify the targeted LDP peer on PE1. |
| (config-router-targeted-peer)# exit | Exit the Router targeted peer mode. |
| (config-router)#exit | Exit the Router mode. |
| (config)#interface xe6 | Specify the interface (xe6) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe6. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe6. |

P1

| | |
|-----------------------------|---|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter the Router mode. |
| (config)# router-id 4.4.4.4 | Configure LDP router ID. |
| (config-router)#exit | Exit the Router mode. |
| (config)#interface xe1/1 | Specify the interface (xe1/1) to be configured. |

| | |
|-----------------------------|---|
| (config-if)#label-switching | Enable label switching on interface xe1/1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe1. |
| (config)#interface xe33 | Specify the interface (xe6) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe33. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe33. |

P2

| | |
|-----------------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter the Router mode. |
| (config)# router-id 5.5.5.5 | Configure LDP router ID. |
| (config-router)#exit | Exit the Router mode. |
| (config)#interface xe6 | Specify the interface (xe6) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe6. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe1. |
| (config)#interface xe47 | Specify the interface (xe47) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe47. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe47. |

Configure VC

Configure the VC. Each VC ID uniquely identifies the Layer-2 circuit among all the Layer-2 circuits.

PE1

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)# mpls l2-circuit VPLS-100-1 10000100 3.3.3.3 | Configure the VC for PE3. In this example, VPLS-100-1 is the VC name, 10000100 is the VC ID, and 3.3.3.3 is the VC endpoint IP address. |
| (config-pseudowire)# mpls l2-circuit VPLS-200-3 30000200 2.2.2.2 | Configure the VC for PE2. In this example, VPLS-200-3 is the VC name, 30000200 is the VC ID, and 2.2.2.2 is the VC endpoint IP address. |

PE2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)# mpls l2-circuit VPLS-200-3 30000200 1.1.1.1 | Configure the VC for PE1. In this example, VPLS-200-3 is the VC name, 30000200 is the VC ID, and 1.1.1.1 is the VC endpoint IP address. |

PE3

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)# mpls l2-circuit VPLS-100-1 10000100 1.1.1.1 mode tagged | Configure the VC for PE1. In this example, VPLS-100-1 is the VC name, 10000100 is the VC ID, and 1.1.1.1 is the VC endpoint IP address. |

Bind Customer Interface to VC**PE1**

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)# mpls l2-circuit VPLS-100-1 10000100 3.3.3.3 | Configure the VC for PE3. In this example, VPLS-100-1 is the VC name, 10000100 is the VC ID, and 3.3.3.3 is the VC endpoint IP address. |
| (config-pseudowire)# mpls l2-circuit VPLS-200-3 30000200 2.2.2.2 | Configure the VC for PE2. In this example, VPLS-200-3 is the VC name, 30000200 is the VC ID, and 2.2.2.2 is the VC endpoint IP address. |
| (config-pseudowire)#exit | Exit pseudowire config mode. |
| (config)#interface xe9.10 switchport | Creates a L2 sub-interface as xe9.10 |
| (config-if)#encapsulation dot1q 10 | Configure the encapsulation as dot1q matching vlan 10 |
| (config-if)#access-if-vpws | Configure access-if-vpws on interface mode |
| (config-acc-if-vpws)#mpls-l2-circuit VPLS-200-3 primary | Configure the VC for PE1 In this example, VPLS-200-3 is the VC name |
| (config-acc-if-vpws)#mpls-l2-circuit VPLS-100-1 secondary | Configure the VC for PE1 In this example, VPLS-200-3 is the VC name |
| (config-acc-if-vpws)#vc-mode revertive | Configured the vc-mode revertive |

PE2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)# mpls l2-circuit VPLS-200-3 30000200 1.1.1.1 | Configure the VC for PE1. In this example, VPLS-200-3 is the VC name, 30000200 is the VC ID, and 1.1.1.1 is the VC endpoint IP address. |
| (config-pseudowire)#exit | Exit pseudowire config mode. |
| (config)#interface xe9.10 switchport | Creates a L2 sub-interface as xe9.10 |
| (config-if)#encapsulation dot1q 10 | Configure the encapsulation as dot1q matching vlan 10 |
| (config-if)#access-if-vpws | Configure access-if-vpws on interface mode |
| (config-acc-if-vpws)#mpls-l2-circuit VPLS-200-3 primary | Configure the VC for PE12 In this example, VPLS-200-3 is the VC name. |

PE3

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)# mpls l2-circuit VPLS-200-3 30000200 1.1.1.1 | Configure the VC for PE1. In this example, VPLS-200-3 is the VC name, 30000200 is the VC ID, and 1.1.1.1 is the VC endpoint IP address. |
| (config-pseudowire)#exit | Exit pseudowire config mode. |
| (config)#interface xe6.100 switchport | Creates a L2 sub-interface as xe6.100 |
| (config-if)#encapsulation dot1q 100 | Configure the encapsulation as dot1q matching vlan 100 |
| (config-if)#access-if-vpws | Configure access-if-vpws on interface mode |
| (config-acc-if-vpws)# mpls-l2-circuit VPLS-100-1 primary | Configure the VC for PE12 In this example, VPLS-100-1 is the VC name |

Validation

```
PE1#sh ldp session      Codes: m - MD5 password is not set/unset.
                        g - GR configuration not set/unset.
                        t - TCP MSS not set/unset.
                        Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | xe6 | Passive | OPERATIONAL | 30 | 02:30:38 |
| | 3.3.3.3 | xe1 | Passive | OPERATIONAL | 30 | 02:30:40 |
| | 5.5.5.5 | xe1 | Passive | OPERATIONAL | 30 | 02:30:38 |
| | 4.4.4.4 | xe6 | Passive | OPERATIONAL | 30 | 02:30:38 |

```
PE1#sh ip ospf neighbor
```

```
Total number of full neighbors: 2
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 5.5.5.5 | 1 | Full/DR | 00:00:30 | 20.20.20.2 | xe1 | 0 |
| 4.4.4.4 | 1 | Full/Backup | 00:00:36 | 40.1.1.2 | xe6 | 0 |

Use the show ldp mpls-l2-circuit (Control Plane) command, and the show mpls vc-table (Forwarding Plane) command, to display complete information about the Layer 2 VC.

If the VC State is UP in the output from the show ldp mpls-l2 circuit command, and the Status is Active in the output of the show mpls vc-table command, a ping from CE1 to CE2 should be successful.

```
PE1#show ldp mpls-l2-circuit
```

| Transport | Client | VC | VC | Local | Remote | Destination |
|-----------|---------|-------|---------------|----------|----------|-------------|
| VC ID | Binding | State | Type | VC Label | VC Label | Address |
| 30000200 | xe9.10 | UP | Ethernet VLAN | 26881 | 26880 | 2.2.2.2 |
| 10000100 | xe9.10 | UP | Ethernet VLAN | 26880 | 26880 | 3.3.3.3 |

```
PE1#sh mpls vc-table
```

```
(m) - Service mapped over multipath transport
```

```
(e) - Service mapped over LDP ECMP
```

| VC-ID | Vlan-ID | Inner-Vlan-ID | Access-Intf | Network-Intf | Out Label | Tunnel-Label | Nextthop | Status | UpTime |
|----------|---------|---------------|-------------|--------------|-----------|--------------|----------|--------|----------|
| 10000100 | N/A | N/A | xe9.10 | - | - | - | - | - | - |
| | 26880 | N/A | 3.3.3.3 | Inactive | - | - | - | - | - |
| 30000200 | N/A | N/A | xe9.10 | xe6 | 26880 | 52480 | 2.2.2.2 | Active | 02:35:05 |

```
PE2#sh ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | xe33 | Active | OPERATIONAL | 30 | 02:37:16 |
| | 4.4.4.4 | xe33 | Passive | OPERATIONAL | 30 | 02:43:19 |

```
PE2#sh ip ospf neighbor
```

```
Total number of full neighbors: 1
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|----------|-----------|-------------|
| 4.4.4.4 | 1 | Full/DR | 00:00:38 | 50.1.1.2 | xe33 | 0 |

```
PE2#sh mpls vc-table
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VC-ID | Vlan-ID | Inner-Vlan-ID | Access-Intf | Network-Intf | Out Label | Tunnel- |
|----------|----------|---------------|-------------|--------------|-----------|---------|
| Label | Nexthop | Status | UpTime | | | |
| 30000200 | N/A | N/A | xe9.10 | xe33 | 26881 | 52481 |
| Active | 02:37:30 | | | | | 1.1.1.1 |

```
PE3#sh ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | xe6 | Active | OPERATIONAL | 30 | 02:38:00 |
| | 5.5.5.5 | xe6 | Passive | OPERATIONAL | 30 | 02:43:52 |

```
PE3#sh ip ospf neighbor
```

```
Total number of full neighbors: 1
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|----------|-----------|-------------|
| 5.5.5.5 | 1 | Full/DR | 00:00:36 | 70.1.1.1 | xe6 | 0 |

```
PE3#sh mpls vc-table
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VC-ID | Vlan-ID | Inner-Vlan-ID | Access-Intf | Network-Intf | Out Label | Tunnel- |
|----------|---------|---------------|-------------|--------------|-----------|---------|
| Label | Nexthop | Status | UpTime | | | |
| 10000100 | N/A | N/A | xe6.100 | - | | |
| | 26880 | N/A | 1.1.1.1 | Inactive | - | |

```
P1#sh ip ospf neighbor
```

```
Total number of full neighbors: 2
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 1.1.1.1 | 1 | Full/DR | 00:00:32 | 40.1.1.1 | xe1/1 | 0 |
| 2.2.2.2 | 1 | Full/Backup | 00:00:31 | 50.1.1.1 | xe33 | 0 |

```
P1#sh ldp session
```

```
Codes: m - MD5 password is not set/unset.
```



```
g - GR configuration not set/unset.  
t - TCP MSS not set/unset.  
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | xe1/1 | Active | OPERATIONAL | 30 | 02:40:25 |
| | 2.2.2.2 | xe33 | Active | OPERATIONAL | 30 | 02:46:28 |

```
P2#sh ldp session
```

```
Codes: m - MD5 password is not set/unset.  
g - GR configuration not set/unset.  
t - TCP MSS not set/unset.  
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | xe47 | Active | OPERATIONAL | 30 | 02:40:56 |
| | 3.3.3.3 | xe6 | Active | OPERATIONAL | 30 | 02:46:49 |

```
P2#sh ip ospf neighbor
```

```
Total number of full neighbors: 2
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 1.1.1.1 | 1 | Full/Backup | 00:00:35 | 20.20.20.1 | xe47 | 0 |
| 3.3.3.3 | 1 | Full/Backup | 00:00:32 | 70.1.1.2 | xe6 | 0 |

```
P2#(config-if)#mpls-l2-circuit t2 service- template ST1
```

L2VPN-L3VPN Stitching

Overview

The L2VPN-L3VPN stitching feature enables seamless interworking between Layer 2 VPNs (VPLS) and the Layer 3 VPN/IP domain within a single OcNOS device. It allows traffic from an L2VPN to be terminated and forwarded into an L3VPN or IP network and vice versa, without requiring physical cabling between devices or interfaces.



Note: When stitching from L3 to L2VPN and host is connected to local ac-subinterface instead of over PW (i.e. local routing), subinterface stats will not be incremented and only parent interface stats will be incremented.



Note: MTU configured under BVI interface which is attached to VPLS should have equal or greater value than maximum MTU configured on PWs. For example, if VPLS has PW1 (MTU 1500) and PW2 (2000) then MTU configured on BVI should be at least 2000 or greater.

Characteristics

- Bridged Virtual Interface (BVI) Based Stitching
 - Uses a logical BVI interface to bridge Layer 2 (VPLS) and Layer 3 (IP/MPLS) domains.
 - BVI functions as the gateway for Layer 3 routing.
- Software-Only Implementation
 - Eliminates the need for external loopback cables or physical interconnections between interfaces.
 - Simplifies network design and reduces hardware dependency.
- MAC-Based Traffic Handling
 - Incoming VPLS traffic is forwarded to Layer 3 if the destination MAC matches the BVI MAC.
 - Outgoing traffic uses ARP resolution to reach hosts in the VPLS domain.

Benefits

- Reduces hardware and operational complexity.
- Saves cost by avoiding extra cables and devices.
- Provides clean integration between L2 and L3 services.
- Enhances service agility and simplifies access-to-core handoff.

Limitation

- Only Q1, Q2, and J2C platforms support this feature.
- The BVI interface must be configured with either untagged encapsulation or single VLAN tagging (only 0x8100 is supported).
- Double-tag (Q-in-Q) encapsulation is not supported on BVI interfaces.

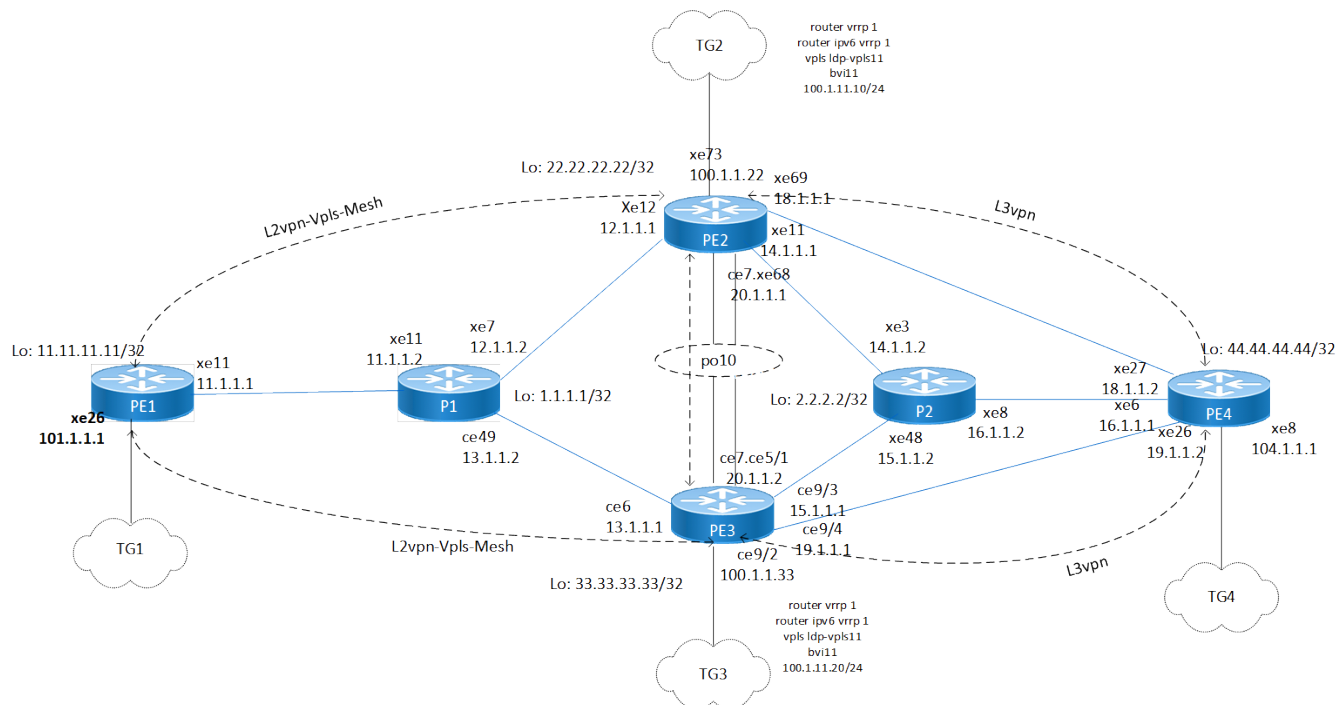
- VLAN rewrite operations are not applicable on BVI interfaces.
- Traffic shaping and queuing features are not supported on BVI interfaces.
- A BVI interface can be associated with only one VPLS instance at a time.
- BVI binding to EVPN is not a valid configuration, as EVPN uses IRB (Integrated Routing and Bridging) for similar functionality.
- Statistics on BVI interface is not supported. User need to check VPLS sub-interface stats or PW stats to see the packet count.
- The configuration of `evpn mpls irb` and `mpls-l2vpn stitching` together is not supported.
- Q1
 - L3 ACL is not supported due to limitation in parsing packet fields after MPLS header.
 - QoS (shaping, queuing, marking, and policing) is not supported.
- Q2
 - For tunnel-terminated packets, ACL matching is performed based on the native packet header.
 - VPLS over BGP-LU with LDP or RSVP is not supported, as the required FEC for BGP-LU is not available.
 - Only Layer 3 ACLs are supported on BVI interfaces.

Configuration

Configure various nodes within the topology to set up a L2VPN L3VPN Stitching session.

Topology

This sample topology provides an L2VPN to IP/L3VPN stitching setup and vice-versa.



PE1

1. Configure LDP:

```
PE1(config)#router ldp
PE1(config-router)#explicit-null
PE1(config-router)#pw-status-tlv
PE1(config-router)#targeted-peer ipv4 22.22.22.22
PE1(config-router-targeted-peer)#exit-targeted-peer-mode
PE1(config-router)#targeted-peer ipv4 33.33.33.33
PE1(config-router-targeted-peer)#exit-targeted-peer-mode
PE1(config-router)#transport-address ipv4 11.11.11.11
PE1(config-router)#commit
PE1(config-router)#exit
```

2. Configure loopback interface:

```
PE1(config)#interface lo
PE1(config-if)#ip address 11.11.11.11/32 secondary
PE1(config-if)#commit
PE1(config-if)#exit
```

3. Configure interface:

```
PE1(config)#interface xe11
PE1(config-if)#description connected-to-pl
PE1(config-if)#load-interval 30
PE1(config-if)#ip address 11.1.1.1/24
PE1(config-if)#mtu 9216
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)#enable-rsvp
PE1(config-if)#commit
PE1(config-if)#exit
```

4. Configure VPLS access interface:

```
PE1(config)#interface xe26.11 switchport
PE1(config-if)#description ldp-vpls
PE1(config-if)#encapsulation dot1q 11
PE1(config-if)#load-interval 30
PE1(config-if)#mtu 9216
PE1(config-if)#access-if-vpls
PE1(config-access-if)#mpls-vpls ldp-vpls11
PE1(config-access-if)#exit
PE1(config-if)#commit
PE1(config-if)#exit
```

5. Configure OSPF

```
PE1(config)#router ospf 100
PE1(config-router)#ospf router-id 11.11.11.11
PE1(config-router)#bfd all-interfaces
PE1(config-router)#network 11.1.1.0/24 area 0.0.0.0
PE1(config-router)#network 11.11.11.11/32 area 0.0.0.0
PE1(config-router)#commit
PE1(config-router)#exit
```

6. Configure BGP:

```
PE1(config)#router bgp 123456
PE1(config-router)#bgp router-id 11.11.11.11
PE1(config-router)#neighbor PG-L2VPN peer-group
PE1(config-router)#neighbor PG-L2VPN remote-as 123456
PE1(config-router)#neighbor PG-L2VPN update-source lo
PE1(config-router)#neighbor PG-L2VPN fall-over bfd multihop
PE1(config-router)#neighbor 22.22.22.22 peer-group PG-L2VPN
PE1(config-router)#neighbor 33.33.33.33 peer-group PG-L2VPN
PE1(config-router)#address-family l2vpn vpls
```

```
PE1(config-router-af)#neighbor PG-L2VPN activate
PE1(config-router-af)#exit-address-family
PE1(config-router)#commit
PE1(config-router)#exit
```

P1

1. Configure interfaces.

```
P1(config)#interface ce49
P1(config-if)#description connected-to-pe3
P1(config-if)#load-interval 30
P1(config-if)#ip address 13.1.1.2/24
P1(config-if)#mtu 9216
P1(config-if)#label-switching
P1(config-if)#enable-ldp ipv4
P1(config-if)#commit
P1(config-if)#exit

P1(config)#interface xe7
P1(config-if)#description connected-to-pe2
P1(config-if)#load-interval 30
P1(config-if)#ip address 12.1.1.2/24
P1(config-if)#mtu 9216
P1(config-if)#label-switching
P1(config-if)#enable-ldp ipv4
P1(config-if)#commit
P1(config-if)#exit

P1(config)#interface xe11
P1(config-if)#description connected-to-pe1
P1(config-if)#load-interval 30
P1(config-if)#ip address 11.1.1.2/24
P1(config-if)#mtu 9216
P1(config-if)#label-switching
P1(config-if)#enable-ldp ipv4
P1(config-if)#commit
P1(config-if)#exit
```

2. Configure loopback interface

```
P1(config)#interface lo
P1(config-if)#ip address 1.1.1.1/32 secondary
P1(config-if)#commit
P1(config-if)#exit
```

3. Configure OSPF routing

```
P1(config)#router ospf 100
P1(config-router)#ospf router-id 1.1.1.1
P1(config-router)#bfd all-interfaces
P1(config-router)#network 1.1.1.1/32 area 0.0.0.0
P1(config-router)#network 11.1.1.0/24 area 0.0.0.0
P1(config-router)#network 12.1.1.0/24 area 0.0.0.0
P1(config-router)#network 13.1.1.0/24 area 0.0.0.0
P1(config-router)#commit
P1(config-router)#exit
```

PE2

1. Enable L2VPN Stitching:

```
PE2(config)#mpls l2vpn-stitching
```



Note: Make sure to enable the command mentioned above for the stitching functionality to operate correctly.

2. Configure the VPLS instance:

```
PE2(config)#mpls vpls ldp-vpls11 11
PE2(config-vpls)#vpls-mtu 1500
PE2(config-vpls-spoke)#vpls-vc vc11
PE2(config-vpls-spoke)#exit-spoke
PE2(config-vpls)#control-word
PE2(config-vpls)#flow-label both
PE2(config-vpls)#signaling ldp
PE2(config-vpls-sig)#vpls-peer 11.11.11.11 vpls-type vlan
PE2(config-vpls-sig)#vpls-peer 33.33.33.33 vpls-type ethernet
PE2(config-vpls-sig)#exit-signaling
PE2(config-vpls)#commit
PE2(config-vpls)#exit-vpls
```

3. Create the L3VPN VRF:

```
PE2(config)#ip vrf l3vpn_vrf11
PE2(config-vrf)#rd 22.22.22.22:11
PE2(config-vrf)#route-target both 11:11
PE2(config-vrf)#commit
PE2(config-vrf)#exit
```

4. Configure LDP:

```
PE2(config)#router ldp
PE2(config-router)#explicit-null
PE2(config-router)#pw-status-tlv
PE2(config-router)#targeted-peer ipv4 11.11.11.11
PE2(config-router-targeted-peer)#exit-targeted-peer-mode
PE2(config-router)#targeted-peer ipv4 33.33.33.33
PE2(config-router-targeted-peer)#exit-targeted-peer-mode
PE2(config-router)#transport-address ipv4 22.22.22.22
PE2(config-router)#commit
PE2(config-router)#exit
```

5. Configure BVI interface (stitching point):

```
PE2(config)#interface bvi11
PE2(config-if)#encapsulation dot1q 11
PE2(config-if)#ip vrf forwarding l3vpn_vrf11
PE2(config-if)#ip address 100.1.11.10/24
PE2(config-if)#ipv6 address 1011::10/64
PE2(config-if)#mtu 9216
PE2(config-if)#access-if-vpls
PE2(config-access-if)#mpls-vpls ldp-vpls11
PE2(config-access-if)#exit
PE2(config-if)#commit
PE2(config-if)#exit
```



Note: Binding VRF and VPLS instance under BVI interface, to ensure that L2vpn to L3vpn stitching and vice-Versa is happening.

6. Configure physical interfaces with MPLS, LDP, and IGP:

```
PE2(config)#interface ce7
PE2(config-if)#channel-group 10 mode active
PE2(config-if)#commit
PE2(config-if)#exit
```

```
PE2(config)#interface xe68
PE2(config-if)#channel-group 10 mode active
PE2(config-if)#commit
PE2(config-if)#exit

PE2(config)#interface po10
PE2(config-if)#description connected-to-pe3
PE2(config-if)#load-interval 30
PE2(config-if)#ip address 20.1.1.1/24
PE2(config-if)#mtu 9216
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)#commit
PE2(config-if)#exit

PE2(config)#interface xe11
PE2(config-if)#description connected-to-p2-rr
PE2(config-if)#load-interval 30
PE2(config-if)#ip address 14.1.1.1/24
PE2(config-if)#mtu 9216
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)#commit
PE2(config-if)#exit

PE2(config)#interface xe12
PE2(config-if)#description connected-to-p
PE2(config-if)#speed 10g
PE2(config-if)#load-interval 30
PE2(config-if)#ip address 12.1.1.1/24
PE2(config-if)#mtu 9216
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)#commit
PE2(config-if)#exit

PE2(config)#interface xe69
PE2(config-if)#description connected-to-pe4
PE2(config-if)#load-interval 30
PE2(config-if)#ip address 18.1.1.1/24
PE2(config-if)#mtu 9216
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)#commit
PE2(config-if)#exit

PE2(config)#interface xe70
PE2(config-if)#speed 1g
PE2(config-if)#load-interval 30
PE2(config-if)#ip address 22.1.1.1/24
PE2(config-if)#mtu 9216
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)#commit
PE2(config-if)#exit
```

7. Configure VPLS access interface:

```
PE2(config)#interface xe73.11 switchport
PE2(config-if)#encapsulation dot1q 11
PE2(config-if)#load-interval 30
PE2(config-if)#mtu 9216
PE2(config-if)#access-if-vpls
PE2(config-access-if)#mpls-vpls ldp-vpls11
PE2(config-if)#commit
PE2(config-if)#exit
```

8. Configure loopback interface:

```
PE2(config)#interface lo
PE2(config-if)#ip address 22.22.22.22/32 secondary
PE2(config-if)#commit
PE2(config-if)#exit
```

9. Configure OSPF IGP:

```
PE2(config)#router ospf 100
PE2(config-router)#ospf router-id 22.22.22.22
PE2(config-router)#bfd all-interfaces
PE2(config-router)#network 12.1.1.0/24 area 0.0.0.0
PE2(config-router)#network 14.1.1.0/24 area 0.0.0.0
PE2(config-router)#network 18.1.1.0/24 area 0.0.0.0
PE2(config-router)#network 20.1.1.0/24 area 0.0.0.0
PE2(config-router)#network 22.1.1.0/24 area 0.0.0.0
PE2(config-router)#network 22.22.22.22/32 area 0.0.0.0
PE2(config-router)#commit
PE2(config-router)#exit
```

10 Configure BGP with VPLS and VPNv4/VPNv6/6PE neighbors:

```
PE2(config)#router bgp 123456
PE2(config-router)#bgp router-id 22.22.22.22
PE2(config-router)#neighbor PG peer-group
PE2(config-router)#neighbor PG-L2VPN peer-group
PE2(config-router)#neighbor PG-L2VPN remote-as 123456
PE2(config-router)#neighbor PG-L2VPN update-source lo
PE2(config-router)#neighbor PG-L2VPN fall-over bfd multihop
PE2(config-router)#neighbor PG-VPNv4 peer-group
PE2(config-router)#neighbor PG-VPNv4 remote-as 123456
PE2(config-router)#neighbor PG-VPNv4 update-source lo
PE2(config-router)#neighbor PG-VPNv4 fall-over bfd multihop
PE2(config-router)#neighbor 2.2.2.2 peer-group PG-VPNv4
PE2(config-router)#neighbor 11.11.11.11 peer-group PG-L2VPN
PE2(config-router)#neighbor 33.33.33.33 peer-group PG-L2VPN
PE2(config-router)#neighbor 641::11 remote-as 100
PE2(config-router)#neighbor 642::11 remote-as 100
PE2(config-router)#neighbor 643::11 remote-as 100
PE2(config-router)#neighbor 644::11 remote-as 100
PE2(config-router)#neighbor 645::11 remote-as 100
PE2(config-router)#neighbor 646::11 remote-as 100
PE2(config-router)#neighbor 647::11 remote-as 100
PE2(config-router)#neighbor 648::11 remote-as 100
PE2(config-router)#neighbor 649::11 remote-as 100
PE2(config-router)#neighbor 650::11 remote-as 100
PE2(config-router-af)#address-family vpnv4 unicast
PE2(config-router-af)#neighbor PG-VPNv4 activate
PE2(config-router-af)#exit-address-family
PE2(config-router)#commit
PE2(config-router)#exit

PE2(config-router-af)#address-family l2vpn vpls
PE2(config-router-af)#neighbor PG-L2VPN activate
PE2(config-router-af)#exit-address-family
PE2(config-router)#commit
PE2(config-router)#exit

PE2(config-router-af)#address-family vpnv6 unicast
PE2(config-router-af)#neighbor PG-VPNv4 activate
PE2(config-router-af)#exit-address-family
PE2(config-router)#commit
PE2(config-router)#exit

PE2(config-router-af)#address-family ipv6 unicast
PE2(config-router-af)#redistribute connected
PE2(config-router-af)#neighbor 641::11 activate
PE2(config-router-af)#neighbor 642::11 activate
PE2(config-router-af)#neighbor 643::11 activate
```



```

PE2(config-router-af)#neighbor 644::11 activate
PE2(config-router-af)#neighbor 645::11 activate
PE2(config-router-af)#neighbor 646::11 activate
PE2(config-router-af)#neighbor 647::11 activate
PE2(config-router-af)#neighbor 648::11 activate
PE2(config-router-af)#neighbor 649::11 activate
PE2(config-router-af)#neighbor 650::11 activate
PE2(config-router-af)#exit-address-family
PE2(config-router)#commit
PE2(config-router)#exit

PE2(config-router-af)#address-family ipv6 labeled-unicast
PE2(config-router-af)#neighbor PG-VPNv4 activate
PE2(config-router-af)#exit-address-family
PE2(config-router)#commit
PE2(config-router)#exit

PE2(config-router-af)#address-family ipv4 vrf l3vpn_vrf11
PE2(config-router-af)#redistribute connected
PE2(config-router-af)#exit-address-family
PE2(config-router)#commit
PE2(config-router)#exit

PE2(config-router-af)#address-family ipv6 vrf l3vpn_vrf11
PE2(config-router-af)#redistribute connected
PE2(config-router-af)#exit-address-family
PE2(config-router)#commit
PE2(config-router)#exit

```

11 Configure VRRPv4/VRRPv6 for gateway redundancy:

```

PE2(config)#router vrrp 1 bvi11
PE2(config-router)#virtual-ip 100.1.11.100
PE2(config-router)#priority 200
PE2(config-router)#v2-compatible
PE2(config-router)#authentication text vrrp
PE2(config-router)#enable
PE2(config-router)#commit
PE2(config-router)#exit

PE2(config)#router ipv6 vrrp 1 bvi11
PE2(config-router)#virtual-ipv6 fe80::11:0:0:1
PE2(config-router)#virtual-ipv6 1011::100
PE2(config-router)#priority 200
PE2(config-router)#enable
PE2(config-router)#commit
PE2(config-router)#exit

```

PE3

1. Enable L2VPN Stitching:

```
PE3(config)#mpls l2vpn-stitching
```



Note: Make sure to enable the command mentioned above for the stitching functionality to operate correctly.

2. Configure the VPLS Instance:

```

PE3(config)#mpls vpls ldp-vpls11 11
PE3(config-vpls)#vpls-mtu 1500
PE3(config-vpls-spoke)#vpls-vc vc111
PE3(config-vpls-spoke)#exit-spoke
PE3(config-vpls)#control-word

```

```

PE3(config-vpls)#flow-label both
PE3(config-vpls)#signaling ldp
PE3(config-vpls-sig)#vpls-peer 11.11.11.11 vpls-type ethernet
PE3(config-vpls-sig)#vpls-peer 22.22.22.22 vpls-type ethernet
PE3(config-vpls-sig)#exit-signaling
PE3(config-vpls)#commit
PE3(config-vpls)#exit-vpls

```

3. Create and Configure the L3VPN VRF:

```

PE3(config)#ip vrf l3vpn_vrf11
PE3(config-vrf)#rd 33.33.33.33:11
PE3(config-vrf)#route-target both 11:11
PE3(config-vrf)#commit
PE3(config-vrf)#exit

```

4. Configure LDP:

```

PE3(config)#router ldp
PE3(config-router)#explicit-null
PE3(config-router)#pw-status-tlv
PE3(config-router)#targeted-peer ipv4 11.11.11.11
PE3(config-router-targeted-peer)#exit-targeted-peer-mode
PE3(config-router)#targeted-peer ipv4 22.22.22.22
PE3(config-router-targeted-peer)#exit-targeted-peer-mode
PE3(config-router)#targeted-peer ipv4 55.55.55.55
PE3(config-router-targeted-peer)#exit-targeted-peer-mode
PE3(config-router)#transport-address ipv4 33.33.33.33
PE3(config-router)#commit
PE3(config-router)#exit

```

5. Configure the BVI Interface (Stitching Point):

```

PE3(config-if)#interface bvi11
PE3(config-if)#encapsulation dot1q 11
PE3(config-if)#ip vrf forwarding l3vpn_vrf11
PE3(config-if)#ip address 100.1.11.20/24
PE3(config-if)#ipv6 address 1011::20/64
PE3(config-if)#mtu 9216
PE3(config-if)#access-if-vpls
PE3(config-access-if)#mpls-vpls ldp-vpls11
PE3(config-access-if)#exit
PE3(config-if)#commit
PE3(config-if)#exit

```



Note: Binding VRF and VPLS instance under BVI interface, to ensure that L2vpn to L3vpn stitching and vice-versa is happening.

6. Configure physical interfaces with MPLS, LDP, and IGP:

```

PE3(config)#interface po10
PE3(config-if)#description connected-to-pe2
PE3(config-if)#load-interval 30
PE3(config-if)#ip address 20.1.1.2/24
PE3(config-if)#mtu 9216
PE3(config-if)#label-switching
PE3(config-if)#enable-ldp ipv4
PE3(config-if)#commit
PE3(config-if)#exit

PE3(config)#interface ce6
PE3(config-if)#description connected-to-p1
PE3(config-if)#load-interval 30
PE3(config-if)#ip address 13.1.1.1/24
PE3(config-if)#mtu 9216

```

```
PE3(config-if)#label-switching
PE3(config-if)#enable-ldp ipv4
PE3(config-if)#commit
PE3(config-if)#exit

PE3(config)#interface ce7
PE3(config-if)#channel-group 10 mode active
PE3(config-if)#commit
PE3(config-if)#exit

PE3(config)#interface ce9/1
PE3(config-if)#channel-group 10 mode active
PE3(config-if)#commit
PE3(config-if)#exit

PE3(config)#interface ce9/2
PE3(config-if)#commit
PE3(config-if)#exit

PE3(config)#interface ce9/3
PE3(config-if)#description connected-to-RR
PE3(config-if)#load-interval 30
PE3(config-if)#ip address 15.1.1.1/24
PE3(config-if)#mtu 9216
PE3(config-if)#label-switching
PE3(config-if)#enable-ldp ipv4
PE3(config-if)#commit
PE3(config-if)#exit

PE3(config)#interface ce9/4
PE3(config-if)#description connected-to-pe4
PE3(config-if)#load-interval 30
PE3(config-if)#ip address 19.1.1.1/24
PE3(config-if)#mtu 9216
PE3(config-if)#label-switching
PE3(config-if)#enable-ldp ipv4
PE3(config-if)#commit
PE3(config-if)#exit

PE3(config)#interface xe3
PE3(config-if)#description connected-to-pe5-spoke
PE3(config-if)#load-interval 30
PE3(config-if)#ip address 21.1.1.1/24
PE3(config-if)#mtu 9216
PE3(config-if)#label-switching
PE3(config-if)#enable-ldp ipv4
PE3(config-if)#commit
PE3(config-if)#exit
```

7. Configure VPLS Access Subinterface:

```
PE3(config)#interface ce9/2.11 switchport
PE3(config-if)#encapsulation dot1q 11
PE3(config-if)#load-interval 30
PE3(config-if)#mtu 9216
PE3(config-if)#access-if-vpls
PE3(config-access-if)#mpls-vpls ldp-vpls11
PE3(config-access-if)#exit
PE3(config-if)#commit
PE3(config-if)#exit
```

8. Configure Loopback Interface:

```
PE3(config)#interface lo
PE3(config-if)#ip address 33.33.33.33/32 secondary
PE3(config-if)#commit
PE3(config-if)#exit
```

9. Configure IGP - OSPF:

```
PE3(config)#router ospf 100
PE3(config-router)#ospf router-id 33.33.33.33
PE3(config-router)#bfd all-interfaces
PE3(config-router)#network 13.1.1.0/24 area 0.0.0.0
PE3(config-router)#network 15.1.1.0/24 area 0.0.0.0
PE3(config-router)#network 19.1.1.0/24 area 0.0.0.0
PE3(config-router)#network 20.1.1.0/24 area 0.0.0.0
PE3(config-router)#network 21.1.1.0/24 area 0.0.0.0
PE3(config-router)#network 33.33.33.33/32 area 0.0.0.0
PE3(config-router)#commit
PE3(config-router)#exit
```

10 Configure BGP with VPLS and VPNv4/VPNv6/6PE neighbors:

```
PE3(config)#router bgp 123456
PE3(config-router)#bgp router-id 33.33.33.33
PE3(config-router)#neighbor PG peer-group
PE3(config-router)#neighbor PG-L2VPN peer-group
PE3(config-router)#neighbor PG-L2VPN remote-as 123456
PE3(config-router)#neighbor PG-L2VPN update-source lo
PE3(config-router)#neighbor PG-L2VPN fall-over bfd multihop
PE3(config-router)#neighbor PG-VPNv4 peer-group
PE3(config-router)#neighbor PG-VPNv4 remote-as 123456
PE3(config-router)#neighbor PG-VPNv4 update-source lo
PE3(config-router)#neighbor PG-VPNv4 fall-over bfd multihop
PE3(config-router)#neighbor 2.2.2.2 peer-group PG-VPNv4
PE3(config-router)#neighbor 11.11.11.11 peer-group PG-L2VPN
PE3(config-router)#neighbor 22.22.22.22 peer-group PG-L2VPN

PE3(config-router)#address-family vpnv4 unicast
PE3(config-router-af)#neighbor PG-VPNv4 activate
PE3(config-router-af)#exit-address-family
PE3(config-router)#commit
PE3(config-router)#exit

PE3(config-router)#address-family l2vpn vpls
PE3(config-router-af)#neighbor PG-L2VPN activate
PE3(config-router-af)#exit-address-family
PE3(config-router)#commit
PE3(config-router)#exit

PE3(config-router)#address-family vpnv6 unicast
PE3(config-router-af)#neighbor PG-VPNv4 activate
PE3(config-router-af)#exit-address-family
PE3(config-router)#commit
PE3(config-router)#exit

PE3(config-router)#address-family ipv6 unicast
PE3(config-router-af)#redistribute connected
PE3(config-router-af)#exit-address-family
PE3(config-router)#commit
PE3(config-router)#exit

PE3(config-router)#address-family ipv6 labeled-unicast
PE3(config-router-af)#neighbor PG-VPNv4 activate
PE3(config-router-af)#exit-address-family
PE3(config-router)#commit
PE3(config-router)#exit

PE3(config-router)#address-family ipv4 vrf l3vpn_vrf11
PE3(config-router-af)#redistribute connected
PE3(config-router-af)#exit-address-family
PE3(config-router)#commit
PE3(config-router)#exit

PE3(config-router)#address-family ipv6 vrf l3vpn_vrf11
PE3(config-router-af)#redistribute connected
PE3(config-router-af)#exit-address-family
```

```
PE3(config-router)#commit
PE3(config-router)#exit
```

11. Configure VRRPv4/VRRPv6 for gateway redundancy:

```
PE3(config)#router vrrp 1 bvi11
PE3(config-router)#virtual-ip 100.1.11.100
PE3(config-router)#priority 150
PE3(config-router)#v2-compatible
PE3(config-router)#authentication text vrrp
PE3(config-router)#enable
PE3(config-router)#commit
PE3(config-router)#exit

PE3(config)#router ipv6 vrrp 1 bvi11
PE3(config-router)#virtual-ipv6 fe80::11:0:0:1
PE3(config-router)#virtual-ipv6 1011::100
PE3(config-router)#priority 150
PE3(config-router)#enable
PE3(config-router)#commit
PE3(config-router)#exit
```

P2

1. Enable LDP:

```
P2(config)#router ldp
```

2. Configure Loopback Interface:

```
P2(config)#interface lo
P2(config-if)#ip address 2.2.2.2/32 secondary
P2(config-if)#commit
P2(config-if)#exit
```

3. Configure Physical Interfaces with IGP and LDP:

```
P2(config)#interface xe3
P2(config-if)#description connected-to-pe2
P2(config-if)#load-interval 30
P2(config-if)#ip address 14.1.1.2/24
P2(config-if)#mtu 9216
P2(config-if)#label-switching
P2(config-if)#enable-ldp ipv4
P2(config-if)#commit
P2(config-if)#exit
```

```
P2(config)#interface xe8
P2(config-if)#description connected-to-pe4
P2(config-if)#load-interval 30
P2(config-if)#ip address 16.1.1.2/24
P2(config-if)#mtu 9216
P2(config-if)#label-switching
P2(config-if)#enable-ldp ipv4
P2(config-if)#commit
P2(config-if)#exit
```

```
P2(config)#interface xe48
P2(config-if)#description connected-to-pe3
P2(config-if)#load-interval 30
P2(config-if)#ip address 15.1.1.2/24
P2(config-if)#mtu 9216
P2(config-if)#label-switching
P2(config-if)#enable-ldp ipv4
P2(config-if)#commit
P2(config-if)#exit
```

4. Configure OSPF:

```
P2(config)#router ospf 100
P2(config-router)#ospf router-id 2.2.2.2
P2(config-router)#bfd all-interfaces
P2(config-router)#network 2.2.2.2/32 area 0.0.0.0
P2(config-router)#network 14.1.1.0/24 area 0.0.0.0
P2(config-router)#network 15.1.1.0/24 area 0.0.0.0
P2(config-router)#network 16.1.1.0/24 area 0.0.0.0
P2(config-router)#commit
P2(config-router)#exit
```

5. Configure BGP Route Reflector:

```
P2(config)#router bgp 123456
P2(config-router)#bgp router-id 2.2.2.2
P2(config-router)#no bgp inbound-route-filter
P2(config-router)#neighbor 22.22.22.22 remote-as 123456
P2(config-router)#neighbor 22.22.22.22 update-source lo
P2(config-router)#neighbor 22.22.22.22 fall-over bfd multihop
P2(config-router)#neighbor 33.33.33.33 remote-as 123456
P2(config-router)#neighbor 33.33.33.33 update-source lo
P2(config-router)#neighbor 33.33.33.33 fall-over bfd multihop
P2(config-router)#neighbor 44.44.44.44 remote-as 123456
P2(config-router)#neighbor 44.44.44.44 update-source lo
P2(config-router)#neighbor 44.44.44.44 fall-over bfd multihop
```

6. Activate Address Families and RR Clients:

```
P2(config-router)#address-family vpnv4 unicast
P2(config-router-af)#neighbor 22.22.22.22 activate
P2(config-router-af)#neighbor 22.22.22.22 route-reflector-client
P2(config-router-af)#neighbor 22.22.22.22 next-hop-self
P2(config-router-af)#neighbor 33.33.33.33 activate
P2(config-router-af)#neighbor 33.33.33.33 route-reflector-client
P2(config-router-af)#neighbor 33.33.33.33 next-hop-self
P2(config-router-af)#neighbor 44.44.44.44 activate
P2(config-router-af)#neighbor 44.44.44.44 route-reflector-client
P2(config-router-af)#neighbor 44.44.44.44 next-hop-self
P2(config-router-af)#exit-address-family
P2(config-router)#commit
P2(config-router)#exit
```

```
P2(config-router-af)#address-family vpnv6 unicast
P2(config-router-af)#neighbor 22.22.22.22 activate
P2(config-router-af)#neighbor 22.22.22.22 route-reflector-client
P2(config-router-af)#neighbor 22.22.22.22 next-hop-self
P2(config-router-af)#neighbor 33.33.33.33 activate
P2(config-router-af)#neighbor 33.33.33.33 route-reflector-client
P2(config-router-af)#neighbor 33.33.33.33 next-hop-self
P2(config-router-af)#neighbor 44.44.44.44 activate
P2(config-router-af)#neighbor 44.44.44.44 route-reflector-client
P2(config-router-af)#neighbor 44.44.44.44 next-hop-self
P2(config-router-af)#exit-address-family
P2(config-router)#commit
P2(config-router)#exit
```

```
P2(config-router-af)#address-family ipv6 unicast
P2(config-router-af)#neighbor 22.22.22.22 activate
P2(config-router-af)#neighbor 22.22.22.22 route-reflector-client
P2(config-router-af)#neighbor 22.22.22.22 next-hop-self
P2(config-router-af)#neighbor 33.33.33.33 activate
P2(config-router-af)#neighbor 33.33.33.33 route-reflector-client
P2(config-router-af)#neighbor 33.33.33.33 next-hop-self
P2(config-router-af)#neighbor 44.44.44.44 activate
P2(config-router-af)#neighbor 44.44.44.44 route-reflector-client
P2(config-router-af)#neighbor 44.44.44.44 next-hop-self
P2(config-router-af)#exit-address-family
P2(config-router)#commit
P2(config-router)#exit
```

```

P2(config-router-af)#address-family ipv6 labeled-unicast
P2(config-router-af)#neighbor 22.22.22.22 activate
P2(config-router-af)#neighbor 22.22.22.22 route-reflector-client
P2(config-router-af)#neighbor 22.22.22.22 next-hop-self
P2(config-router-af)#neighbor 33.33.33.33 activate
P2(config-router-af)#neighbor 33.33.33.33 route-reflector-client
P2(config-router-af)#neighbor 33.33.33.33 next-hop-self
P2(config-router-af)#neighbor 44.44.44.44 activate
P2(config-router-af)#neighbor 44.44.44.44 route-reflector-client
P2(config-router-af)#neighbor 44.44.44.44 next-hop-self
P2(config-router-af)#exit-address-family
P2(config-router)#commit
P2(config-router)#exit

```

PE4

1. Create VRFs:

```

PE4(config)#ip vrf l3vpn_vrf11
PE4(config-vrf)#rd 44.44.44.44:11
PE4(config-vrf)#route-target both 11:11
PE4(config-vrf)#ip vrf management
PE4(config-vrf)#commit
PE4(config-vrf)#exit

```

2. Enable LDP Globally:

```

PE4(config)#router ldp

```

3. Configure Loopback Interface:

```

PE4(config)#interface lo
PE4(config-if)#ip address 44.44.44.44/32 secondary
P2(config-if)#commit
P2(config-if)#exit

```

4. Configure physical interfaces with MPLS, LDP, and IGP:

```

PE4(config)#interface xe5
PE4(config-if)#description connected-to-pe5
PE4(config-if)#speed 10g
PE4(config-if)#load-interval 30
PE4(config-if)#ip address 17.1.1.1/24
PE4(config-if)#mtu 9216
PE4(config-if)#label-switching
PE4(config-if)#enable-ldp ipv4
PE4(config-if)#enable-rsvp
PE4(config-if)#commit
PE4(config-if)#exit

PE4(config)#interface xe6
PE4(config-if)#description connected-to-p2-rr
PE4(config-if)#speed 10g
PE4(config-if)#load-interval 30
PE4(config-if)#ip address 16.1.1.1/24
PE4(config-if)#mtu 9216
PE4(config-if)#label-switching
PE4(config-if)#enable-ldp ipv4
PE4(config-if)#enable-rsvp
PE4(config-if)#commit
PE4(config-if)#exit

PE4(config)#interface xe26
PE4(config-if)#description connected-to-pe3
PE4(config-if)#load-interval 30
PE4(config-if)#ip address 19.1.1.2/24

```

```
PE4(config-if)#mtu 9216
PE4(config-if)#label-switching
PE4(config-if)#enable-ldp ipv4
PE4(config-if)#enable-rsvp
PE4(config-if)#commit
PE4(config-if)#exit

PE4(config)#interface xe27
PE4(config-if)#description connected-to-pe2
PE4(config-if)#load-interval 30
PE4(config-if)#ip address 18.1.1.2/24
PE4(config-if)#mtu 9216
PE4(config-if)#label-switching
PE4(config-if)#enable-ldp ipv4
PE4(config-if)#enable-rsvp
PE4(config-if)#commit
PE4(config-if)#exit
```

5. Configure Access Interface (VRF-bound):

```
PE4(config)#interface xe8.11
PE4(config-if)#encapsulation dot1q 11
PE4(config-if)#load-interval 30
PE4(config-if)#ip vrf forwarding l3vpn_vrf11
PE4(config-if)#ip address 104.1.11.1/24
PE4(config-if)#ipv6 address 1411::1/64
PE4(config-if)#mtu 9216
```

6. Configure OSPF for Core Routing:

```
PE4(config)#router ospf 100
PE4(config-router)#ospf router-id 44.44.44.44
PE4(config-router)#bfd all-interfaces
PE4(config-router)#network 16.1.1.0/24 area 0.0.0.0
PE4(config-router)#network 17.1.1.0/24 area 0.0.0.0
PE4(config-router)#network 18.1.1.0/24 area 0.0.0.0
PE4(config-router)#network 19.1.1.0/24 area 0.0.0.0
PE4(config-router)#network 44.44.44.44/32 area 0.0.0.0
PE4(config-router)#commit
PE4(config-router)#exit
```

7. Configure BGP for VPNv4/VPNv6/6PE:

```
PE4(config)#router bgp 123456
PE4(config-router)#bgp router-id 44.44.44.44
PE4(config-router)#neighbor 2.2.2.2 remote-as 123456
PE4(config-router)#neighbor 2.2.2.2 update-source lo
PE4(config-router)#neighbor 2.2.2.2 fall-over bfd multihop
PE4(config-router)#neighbor 1641::2 remote-as 100
PE4(config-router)#neighbor 1642::2 remote-as 100
PE4(config-router)#neighbor 1643::2 remote-as 100
PE4(config-router)#neighbor 1644::2 remote-as 100
PE4(config-router)#neighbor 1645::2 remote-as 100
PE4(config-router)#neighbor 1646::2 remote-as 100
PE4(config-router)#neighbor 1647::2 remote-as 100
PE4(config-router)#neighbor 1648::2 remote-as 100
PE4(config-router)#neighbor 1649::2 remote-as 100
PE4(config-router)#neighbor 1650::2 remote-as 100
PE4(config-router)#commit
PE4(config-router)#exit
```

8. Activate BGP Address Families

```
PE4(config-router)#address-family vpnv4 unicast
PE4(config-router-af)#neighbor 2.2.2.2 activate
PE4(config-router-af)#exit-address-family
PE4(config-router)#commit
PE4(config-router)#exit
```



```
PE4(config-router)#address-family vpnv6 unicast
PE4(config-router-af)#neighbor 2.2.2.2 activate
PE4(config-router-af)#exit-address-family
PE4(config-router)#commit
PE4(config-router)#exit

PE4(config-router)#address-family ipv6 unicast
PE4(config-router-af)#redistribute connected
PE4(config-router-af)#neighbor 1641::2 activate
PE4(config-router-af)#neighbor 1642::2 activate
PE4(config-router-af)#neighbor 1643::2 activate
PE4(config-router-af)#neighbor 1644::2 activate
PE4(config-router-af)#neighbor 1645::2 activate
PE4(config-router-af)#neighbor 1646::2 activate
PE4(config-router-af)#neighbor 1647::2 activate
PE4(config-router-af)#neighbor 1648::2 activate
PE4(config-router-af)#neighbor 1649::2 activate
PE4(config-router-af)#neighbor 1650::2 activate
PE4(config-router-af)#exit-address-family
PE4(config-router)#commit
PE4(config-router)#exit

PE4(config-router)#address-family ipv6 labeled-unicast
PE4(config-router-af)#neighbor 2.2.2.2 activate
PE4(config-router-af)#exit-address-family
PE4(config-router)#commit
PE4(config-router)#exit
```

9. VRF IPv4 & IPv6 Routing for Customers

```
PE4(config-router)# address-family ipv4 vrf l3vpn_vrf11
PE4(config-router-af)#redistribute connected
PE4(config-router-af)#neighbor 104.1.11.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.11.2 activate
PE4(config-router-af)#neighbor 104.1.12.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.12.2 activate
PE4(config-router-af)#neighbor 104.1.13.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.13.2 activate
PE4(config-router-af)#neighbor 104.1.14.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.14.2 activate
PE4(config-router-af)#neighbor 104.1.15.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.15.2 activate
PE4(config-router-af)#neighbor 104.1.16.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.16.2 activate
PE4(config-router-af)#neighbor 104.1.17.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.17.2 activate
PE4(config-router-af)#neighbor 104.1.18.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.18.2 activate
PE4(config-router-af)#neighbor 104.1.19.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.19.2 activate
PE4(config-router-af)#neighbor 104.1.20.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.20.2 activate
PE4(config-router-af)#neighbor 104.1.21.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.21.2 activate
PE4(config-router-af)#neighbor 104.1.22.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.22.2 activate
PE4(config-router-af)#neighbor 104.1.23.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.23.2 activate
PE4(config-router-af)#neighbor 104.1.24.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.24.2 activate
PE4(config-router-af)#neighbor 104.1.25.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.25.2 activate
PE4(config-router-af)#neighbor 104.1.26.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.26.2 activate
PE4(config-router-af)#neighbor 104.1.27.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.27.2 activate
PE4(config-router-af)#neighbor 104.1.28.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.28.2 activate
```

```
PE4(config-router-af)#neighbor 104.1.29.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.29.2 activate
PE4(config-router-af)#neighbor 104.1.30.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.30.2 activate
PE4(config-router-af)#neighbor 104.1.31.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.31.2 activate
PE4(config-router-af)#neighbor 104.1.32.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.32.2 activate
PE4(config-router-af)#neighbor 104.1.33.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.33.2 activate
PE4(config-router-af)#neighbor 104.1.34.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.34.2 activate
PE4(config-router-af)#neighbor 104.1.35.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.35.2 activate
PE4(config-router-af)#neighbor 104.1.36.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.36.2 activate
PE4(config-router-af)#neighbor 104.1.37.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.37.2 activate
PE4(config-router-af)#neighbor 104.1.38.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.38.2 activate
PE4(config-router-af)#neighbor 104.1.39.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.39.2 activate
PE4(config-router-af)#neighbor 104.1.40.2 remote-as 100
PE4(config-router-af)#neighbor 104.1.40.2 activate
PE4(config-router-af)#exit-address-family
PE4(config-router)#commit
PE4(config-router)#exit
```

```
PE4(config-router)#address-family ipv6 vrf l3vpn_vrf11
PE4(config-router-af)#redistribute connected
PE4(config-router-af)#neighbor 1411::2 remote-as 100
PE4(config-router-af)#neighbor 1411::2 activate
PE4(config-router-af)#neighbor 1412::2 remote-as 100
PE4(config-router-af)#neighbor 1412::2 activate
PE4(config-router-af)#neighbor 1413::2 remote-as 100
PE4(config-router-af)#neighbor 1413::2 activate
PE4(config-router-af)#neighbor 1414::2 remote-as 100
PE4(config-router-af)#neighbor 1414::2 activate
PE4(config-router-af)#neighbor 1415::2 remote-as 100
PE4(config-router-af)#neighbor 1415::2 activate
PE4(config-router-af)#neighbor 1416::2 remote-as 100
PE4(config-router-af)#neighbor 1416::2 activate
PE4(config-router-af)#neighbor 1417::2 remote-as 100
PE4(config-router-af)#neighbor 1417::2 activate
PE4(config-router-af)#neighbor 1418::2 remote-as 100
PE4(config-router-af)#neighbor 1418::2 activate
PE4(config-router-af)#neighbor 1419::2 remote-as 100
PE4(config-router-af)#neighbor 1419::2 activate
PE4(config-router-af)#neighbor 1420::2 remote-as 100
PE4(config-router-af)#neighbor 1420::2 activate
PE4(config-router-af)#neighbor 1421::2 remote-as 100
PE4(config-router-af)#neighbor 1421::2 activate
PE4(config-router-af)#neighbor 1422::2 remote-as 100
PE4(config-router-af)#neighbor 1422::2 activate
PE4(config-router-af)#neighbor 1423::2 remote-as 100
PE4(config-router-af)#neighbor 1423::2 activate
PE4(config-router-af)#neighbor 1424::2 remote-as 100
PE4(config-router-af)#neighbor 1424::2 activate
PE4(config-router-af)#neighbor 1425::2 remote-as 100
PE4(config-router-af)#neighbor 1425::2 activate
PE4(config-router-af)#neighbor 1426::2 remote-as 100
PE4(config-router-af)#neighbor 1426::2 activate
PE4(config-router-af)#neighbor 1427::2 remote-as 100
PE4(config-router-af)#neighbor 1427::2 activate
PE4(config-router-af)#neighbor 1428::2 remote-as 100
PE4(config-router-af)#neighbor 1428::2 activate
PE4(config-router-af)#neighbor 1429::2 remote-as 100
PE4(config-router-af)#neighbor 1429::2 activate
```

```

PE4(config-router-af)#neighbor 1430::2 remote-as 100
PE4(config-router-af)#neighbor 1430::2 activate
PE4(config-router-af)#neighbor 1431::2 remote-as 100
PE4(config-router-af)#neighbor 1431::2 activate
PE4(config-router-af)#neighbor 1432::2 remote-as 100
PE4(config-router-af)#neighbor 1432::2 activate
PE4(config-router-af)#neighbor 1433::2 remote-as 100
PE4(config-router-af)#neighbor 1433::2 activate
PE4(config-router-af)#neighbor 1434::2 remote-as 100
PE4(config-router-af)#neighbor 1434::2 activate
PE4(config-router-af)#neighbor 1435::2 remote-as 100
PE4(config-router-af)#neighbor 1435::2 activate
PE4(config-router-af)#neighbor 1436::2 remote-as 100
PE4(config-router-af)#neighbor 1436::2 activate
PE4(config-router-af)#neighbor 1437::2 remote-as 100
PE4(config-router-af)#neighbor 1437::2 activate
PE4(config-router-af)#neighbor 1438::2 remote-as 100
PE4(config-router-af)#neighbor 1438::2 activate
PE4(config-router-af)#neighbor 1439::2 remote-as 100
PE4(config-router-af)#neighbor 1439::2 activate
PE4(config-router-af)#neighbor 1440::2 remote-as 100
PE4(config-router-af)#neighbor 1440::2 activate
PE4(config-router-af)#exit-address-family
PE4(config-router)#commit
PE4(config-router)#exit

```

Running Configuration

PE1

```

!
router ldp
  explicit-null
  pw-status-tlv
  targeted-peer ipv4 22.22.22.22
  exit-targeted-peer-mode
  targeted-peer ipv4 33.33.33.33
  exit-targeted-peer-mode
  transport-address ipv4 11.11.11.11
!
interface lo
  ip address 11.11.11.11/32 secondary
!
interface xel1
  description connected-to-pl
  load-interval 30
  ip address 11.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface xe26.11 switchport
  description ldp-vpls
  encapsulation dot1q 11
  load-interval 30
  mtu 9216
  access-if-vpls
  mpls-vpls ldp-vpls11
!
router ospf 100
  ospf router-id 11.11.11.11
  bfd all-interfaces
  network 11.1.1.0/24 area 0.0.0.0
  network 11.11.11.11/32 area 0.0.0.0
!

```

```
router bgp 123456
  bgp router-id 11.11.11.11
  neighbor PG-L2VPN peer-group
  neighbor PG-L2VPN remote-as 123456
  neighbor PG-L2VPN update-source lo
  neighbor PG-L2VPN fall-over bfd multihop
  neighbor 22.22.22.22 peer-group PG-L2VPN
  neighbor 33.33.33.33 peer-group PG-L2VPN
  !
  address-family l2vpn vpls
  neighbor PG-L2VPN activate
  exit-address-family
  !
exit
```

P1

```
!
interface ce49
  description connected-to-pe3
  load-interval 30
  ip address 13.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
!
interface lo
  ip address 1.1.1.1/32 secondary
!
interface xe7
  description connected-to-pe2
  load-interval 30
  ip address 12.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
!
interface xel1
  description connected-to-pe1
  load-interval 30
  ip address 11.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
!
router ospf 100
  ospf router-id 1.1.1.1
  bfd all-interfaces
  network 1.1.1.1/32 area 0.0.0.0
  network 11.1.1.0/24 area 0.0.0.0
  network 12.1.1.0/24 area 0.0.0.0
  network 13.1.1.0/24 area 0.0.0.0
!
```

PE2

```
!
mpls l2vpn-stitching
!
mpls vpls ldp-vpls11 11
  vpls-mtu 1500
  vpls-vc vc11
  exit-spoke
  control-word
  flow-label both
  signaling ldp
  vpls-peer 11.11.11.11 vpls-type vlan
```

```
    vpls-peer 33.33.33.33 vpls-type ethernet
  exit-signaling
exit-vpls
!
hostname PE2
!
ip vrf l3vpn_vrf11
  rd 22.22.22.22:11
  route-target both 11:11
!
router ldp
  explicit-null
  pw-status-tlv
  targeted-peer ipv4 11.11.11.11
    exit-targeted-peer-mode
  targeted-peer ipv4 33.33.33.33
    exit-targeted-peer-mode
  transport-address ipv4 22.22.22.22
!
interface po10
  description connected-to-pe3
  load-interval 30
  ip address 20.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4

!
interface bv11
  encapsulation dot1q 11
  ip vrf forwarding l3vpn_vrf11
  ip address 100.1.11.10/24
  ipv6 address 1011::10/64
  mtu 9216
  access-if-vpls
    mpls-vpls ldp-vpls11
!
interface ce7
  channel-group 10 mode active
!
interface lo
  ip address 22.22.22.22/32 secondary
!
interface xe11
  description connected-to-p2-rr
  load-interval 30
  ip address 14.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4

!
interface xe12
  description connected-to-p
  speed 10g
  load-interval 30
  ip address 12.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4

!
interface xe68
  channel-group 10 mode active
!
interface xe69
  description connected-to-pe4
  load-interval 30
```

```
ip address 18.1.1.1/24
mtu 9216
label-switching
enable-ldp ipv4

!
interface xe70
speed 1g
load-interval 30
ip address 22.1.1.1/24
mtu 9216
label-switching
enable-ldp ipv4

!
interface xe73.11 switchport
encapsulation dot1q 11
load-interval 30
mtu 9216
access-if-vpls
mpls-vpls ldp-vpls11
!
router ospf 100
ospf router-id 22.22.22.22
bfd all-interfaces
network 12.1.1.0/24 area 0.0.0.0
network 14.1.1.0/24 area 0.0.0.0
network 18.1.1.0/24 area 0.0.0.0
network 20.1.1.0/24 area 0.0.0.0
network 22.1.1.0/24 area 0.0.0.0
network 22.22.22.22/32 area 0.0.0.0
!
router bgp 123456
bgp router-id 22.22.22.22
neighbor PG peer-group
neighbor PG-L2VPN peer-group
neighbor PG-L2VPN remote-as 123456
neighbor PG-L2VPN update-source lo
neighbor PG-L2VPN fall-over bfd multihop
neighbor PG-VPNv4 peer-group
neighbor PG-VPNv4 remote-as 123456
neighbor PG-VPNv4 update-source lo
neighbor PG-VPNv4 fall-over bfd multihop
neighbor 2.2.2.2 peer-group PG-VPNv4
neighbor 11.11.11.11 peer-group PG-L2VPN
neighbor 33.33.33.33 peer-group PG-L2VPN
neighbor 641::11 remote-as 100
neighbor 642::11 remote-as 100
neighbor 643::11 remote-as 100
neighbor 644::11 remote-as 100
neighbor 645::11 remote-as 100
neighbor 646::11 remote-as 100
neighbor 647::11 remote-as 100
neighbor 648::11 remote-as 100
neighbor 649::11 remote-as 100
neighbor 650::11 remote-as 100
!
address-family vpnv4 unicast
neighbor PG-VPNv4 activate
exit-address-family
!
address-family l2vpn vpls
neighbor PG-L2VPN activate
exit-address-family
!
address-family vpnv6 unicast
neighbor PG-VPNv4 activate
exit-address-family
```

```

!
address-family ipv6 unicast
redistribute connected
neighbor 641::11 activate
neighbor 642::11 activate
neighbor 643::11 activate
neighbor 644::11 activate
neighbor 645::11 activate
neighbor 646::11 activate
neighbor 647::11 activate
neighbor 648::11 activate
neighbor 649::11 activate
neighbor 650::11 activate
exit-address-family
!
address-family ipv6 labeled-unicast
neighbor PG-VPNv4 activate
exit-address-family
!
address-family ipv4 vrf l3vpn_vrf11
redistribute connected
exit-address-family
!
address-family ipv6 vrf l3vpn_vrf11
redistribute connected
exit-address-family
!
exit
!
router vrrp 1 bvi11
virtual-ip 100.1.11.100
priority 200
v2-compatible
authentication text vrrp
enable
!
router ipv6 vrrp 1 bvi11
virtual-ipv6 fe80::11:0:0:1
virtual-ipv6 1011::100
priority 200
enable
!

```

PE3

```

!
mpls l2vpn-stitching
!
mpls vpls ldp-vpls11 11
vpls-mtu 1500
vpls-vc vc111
exit-spoke
control-word
flow-label both
signaling ldp
vpls-peer 11.11.11.11 vpls-type ethernet
vpls-peer 22.22.22.22 vpls-type ethernet
exit-signaling
exit-vpls
!
hostname PE3
port ce3 breakout 4X10g
port ce9 breakout 4X10g
port ce20 breakout 4X10g
!
ip vrf l3vpn_vrf11
rd 33.33.33.33:11

```

```
route-target both 11:11
!
router ldp
explicit-null
pw-status-tlv
targeted-peer ipv4 11.11.11.11
exit-targeted-peer-mode
targeted-peer ipv4 22.22.22.22
exit-targeted-peer-mode
targeted-peer ipv4 55.55.55.55
exit-targeted-peer-mode
transport-address ipv4 33.33.33.33
!
interface po10
description connected-to-pe2
load-interval 30
ip address 20.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
!
interface bvi11
encapsulation dot1q 11
ip vrf forwarding l3vpn_vrf11
ip address 100.1.11.20/24
ipv6 address 1011::20/64
mtu 9216
access-if-vpls
mpls-vpls ldp-vpls11
!
interface ce6
description connected-to-pl
load-interval 30
ip address 13.1.1.1/24
mtu 9216
label-switching
enable-ldp ipv4
!
interface ce7
channel-group 10 mode active
!
interface ce9/1
channel-group 10 mode active
!
interface ce9/2
!
interface ce9/2.11 switchport
encapsulation dot1q 11
load-interval 30
mtu 9216
access-if-vpls
mpls-vpls ldp-vpls11
!
interface ce9/3
description connected-to-RR
load-interval 30
ip address 15.1.1.1/24
mtu 9216
label-switching
enable-ldp ipv4
!
interface ce9/4
description connected-to-pe4
load-interval 30
ip address 19.1.1.1/24
mtu 9216
label-switching
enable-ldp ipv4
```



```
!  
interface lo  
  ip address 33.33.33.33/32 secondary  
!  
interface xe3  
  description connected-to-pe5-spoke  
  load-interval 30  
  ip address 21.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
!  
router ospf 100  
  ospf router-id 33.33.33.33  
  bfd all-interfaces  
  network 13.1.1.0/24 area 0.0.0.0  
  network 15.1.1.0/24 area 0.0.0.0  
  network 19.1.1.0/24 area 0.0.0.0  
  network 20.1.1.0/24 area 0.0.0.0  
  network 21.1.1.0/24 area 0.0.0.0  
  network 33.33.33.33/32 area 0.0.0.0  
!  
router bgp 123456  
  bgp router-id 33.33.33.33  
  neighbor PG peer-group  
  neighbor PG-L2VPN peer-group  
  neighbor PG-L2VPN remote-as 123456  
  neighbor PG-L2VPN update-source lo  
  neighbor PG-L2VPN fall-over bfd multihop  
  neighbor PG-VPNv4 peer-group  
  neighbor PG-VPNv4 remote-as 123456  
  neighbor PG-VPNv4 update-source lo  
  neighbor PG-VPNv4 fall-over bfd multihop  
  neighbor 2.2.2.2 peer-group PG-VPNv4  
  neighbor 11.11.11.11 peer-group PG-L2VPN  
  neighbor 22.22.22.22 peer-group PG-L2VPN  
!  
  address-family vpnv4 unicast  
  neighbor PG-VPNv4 activate  
  exit-address-family  
!  
  address-family l2vpn vpls  
  neighbor PG-L2VPN activate  
  exit-address-family  
!  
  address-family vpnv6 unicast  
  neighbor PG-VPNv4 activate  
  exit-address-family  
!  
  address-family ipv6 unicast  
  redistribute connected  
  exit-address-family  
!  
  address-family ipv6 labeled-unicast  
  neighbor PG-VPNv4 activate  
  exit-address-family  
!  
  address-family ipv4 vrf l3vpn_vrf11  
  redistribute connected  
  exit-address-family  
!  
  address-family ipv6 vrf l3vpn_vrf11  
  redistribute connected  
  exit-address-family  
!  
exit  
!
```

```
router vrrp 1 bv111
virtual-ip 100.1.11.100
priority 150
v2-compatible
authentication text vrrp
enable
!
router ipv6 vrrp 1 bv111
virtual-ipv6 fe80::11:0:0:1
virtual-ipv6 1011::100
priority 150
enable
!
```

P2

```
!
router ldp
!
interface lo
ip address 2.2.2.2/32 secondary
!
interface xe3
description connected-to-pe2
load-interval 30
ip address 14.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
!
interface xe8
description connected-to-pe4
load-interval 30
ip address 16.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
!
interface xe48
description connected-to-pe3
load-interval 30
ip address 15.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
!
router ospf 100
ospf router-id 2.2.2.2
bfd all-interfaces
network 2.2.2.2/32 area 0.0.0.0
network 14.1.1.0/24 area 0.0.0.0
network 15.1.1.0/24 area 0.0.0.0
network 16.1.1.0/24 area 0.0.0.0
!
router bgp 123456
bgp router-id 2.2.2.2
no bgp inbound-route-filter
neighbor 22.22.22.22 remote-as 123456
neighbor 22.22.22.22 update-source lo
neighbor 22.22.22.22 fall-over bfd multihop
neighbor 33.33.33.33 remote-as 123456
neighbor 33.33.33.33 update-source lo
neighbor 33.33.33.33 fall-over bfd multihop
neighbor 44.44.44.44 remote-as 123456
neighbor 44.44.44.44 update-source lo
neighbor 44.44.44.44 fall-over bfd multihop
!
```

```

address-family vpnv4 unicast
neighbor 22.22.22.22 activate
neighbor 22.22.22.22 route-reflector-client
neighbor 22.22.22.22 next-hop-self
neighbor 33.33.33.33 activate
neighbor 33.33.33.33 route-reflector-client
neighbor 33.33.33.33 next-hop-self
neighbor 44.44.44.44 activate
neighbor 44.44.44.44 route-reflector-client
neighbor 44.44.44.44 next-hop-self
exit-address-family
!
address-family vpnv6 unicast
neighbor 22.22.22.22 activate
neighbor 22.22.22.22 route-reflector-client
neighbor 22.22.22.22 next-hop-self
neighbor 33.33.33.33 activate
neighbor 33.33.33.33 route-reflector-client
neighbor 33.33.33.33 next-hop-self
neighbor 44.44.44.44 activate
neighbor 44.44.44.44 route-reflector-client
neighbor 44.44.44.44 next-hop-self
exit-address-family
!
address-family ipv6 unicast
neighbor 22.22.22.22 activate
neighbor 22.22.22.22 route-reflector-client
neighbor 22.22.22.22 next-hop-self
neighbor 33.33.33.33 activate
neighbor 33.33.33.33 route-reflector-client
neighbor 33.33.33.33 next-hop-self
neighbor 44.44.44.44 activate
neighbor 44.44.44.44 route-reflector-client
neighbor 44.44.44.44 next-hop-self
exit-address-family
!
address-family ipv6 labeled-unicast
neighbor 22.22.22.22 activate
neighbor 22.22.22.22 route-reflector-client
neighbor 22.22.22.22 next-hop-self
neighbor 33.33.33.33 activate
neighbor 33.33.33.33 route-reflector-client
neighbor 33.33.33.33 next-hop-self
neighbor 44.44.44.44 activate
neighbor 44.44.44.44 route-reflector-client
neighbor 44.44.44.44 next-hop-self
exit-address-family
!

```

PE4

```

!
ip vrf l3vpn_vrf11
  rd 44.44.44.44:11
  route-target both 11:11
!
ip vrf management
!
router ldp
!
interface lo
  ip address 44.44.44.44/32 secondary
!
interface xe5
  description connected-to-pe5
  speed 10g
  load-interval 30

```

```
ip address 17.1.1.1/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface xe6
description connected-to-p2-rr
speed 10g
load-interval 30
ip address 16.1.1.1/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface xe8.11
encapsulation dot1q 11
load-interval 30
ip vrf forwarding l3vpn_vrf11
ip address 104.1.11.1/24
ipv6 address 1411::1/64
mtu 9216
!
interface xe26
description connected-to-pe3
load-interval 30
ip address 19.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface xe27
description connectedt-to-pe2
load-interval 30
ip address 18.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
router ospf 100
ospf router-id 44.44.44.44
bfd all-interfaces
network 16.1.1.0/24 area 0.0.0.0
network 17.1.1.0/24 area 0.0.0.0
network 18.1.1.0/24 area 0.0.0.0
network 19.1.1.0/24 area 0.0.0.0
network 44.44.44.44/32 area 0.0.0.0
!
router bgp 123456
bgp router-id 44.44.44.44
neighbor 2.2.2.2 remote-as 123456
neighbor 2.2.2.2 update-source lo
neighbor 2.2.2.2 fall-over bfd multihop
neighbor 1641::2 remote-as 100
neighbor 1642::2 remote-as 100
neighbor 1643::2 remote-as 100
neighbor 1644::2 remote-as 100
neighbor 1645::2 remote-as 100
neighbor 1646::2 remote-as 100
neighbor 1647::2 remote-as 100
neighbor 1648::2 remote-as 100
neighbor 1649::2 remote-as 100
neighbor 1650::2 remote-as 100
!
address-family vpnv4 unicast
```

```
neighbor 2.2.2.2 activate
exit-address-family
!
address-family vpnv6 unicast
neighbor 2.2.2.2 activate
exit-address-family
!
address-family ipv6 unicast
redistribute connected
neighbor 1641::2 activate
neighbor 1642::2 activate
neighbor 1643::2 activate
neighbor 1644::2 activate
neighbor 1645::2 activate
neighbor 1646::2 activate
neighbor 1647::2 activate
neighbor 1648::2 activate
neighbor 1649::2 activate
neighbor 1650::2 activate
exit-address-family
!
address-family ipv6 labeled-unicast
neighbor 2.2.2.2 activate
exit-address-family
!
address-family ipv4 vrf l3vpn_vrf11
redistribute connected
neighbor 104.1.11.2 remote-as 100
neighbor 104.1.11.2 activate
neighbor 104.1.12.2 remote-as 100
neighbor 104.1.12.2 activate
neighbor 104.1.13.2 remote-as 100
neighbor 104.1.13.2 activate
neighbor 104.1.14.2 remote-as 100
neighbor 104.1.14.2 activate
neighbor 104.1.15.2 remote-as 100
neighbor 104.1.15.2 activate
neighbor 104.1.16.2 remote-as 100
neighbor 104.1.16.2 activate
neighbor 104.1.17.2 remote-as 100
neighbor 104.1.17.2 activate
neighbor 104.1.18.2 remote-as 100
neighbor 104.1.18.2 activate
neighbor 104.1.19.2 remote-as 100
neighbor 104.1.19.2 activate
neighbor 104.1.20.2 remote-as 100
neighbor 104.1.20.2 activate
neighbor 104.1.21.2 remote-as 100
neighbor 104.1.21.2 activate
neighbor 104.1.22.2 remote-as 100
neighbor 104.1.22.2 activate
neighbor 104.1.23.2 remote-as 100
neighbor 104.1.23.2 activate
neighbor 104.1.24.2 remote-as 100
neighbor 104.1.24.2 activate
neighbor 104.1.25.2 remote-as 100
neighbor 104.1.25.2 activate
neighbor 104.1.26.2 remote-as 100
neighbor 104.1.26.2 activate
neighbor 104.1.27.2 remote-as 100
neighbor 104.1.27.2 activate
neighbor 104.1.28.2 remote-as 100
neighbor 104.1.28.2 activate
neighbor 104.1.29.2 remote-as 100
neighbor 104.1.29.2 activate
neighbor 104.1.30.2 remote-as 100
neighbor 104.1.30.2 activate
neighbor 104.1.31.2 remote-as 100
```

```
neighbor 104.1.31.2 activate
neighbor 104.1.32.2 remote-as 100
neighbor 104.1.32.2 activate
neighbor 104.1.33.2 remote-as 100
neighbor 104.1.33.2 activate
neighbor 104.1.34.2 remote-as 100
neighbor 104.1.34.2 activate
neighbor 104.1.35.2 remote-as 100
neighbor 104.1.35.2 activate
neighbor 104.1.36.2 remote-as 100
neighbor 104.1.36.2 activate
neighbor 104.1.37.2 remote-as 100
neighbor 104.1.37.2 activate
neighbor 104.1.38.2 remote-as 100
neighbor 104.1.38.2 activate
neighbor 104.1.39.2 remote-as 100
neighbor 104.1.39.2 activate
neighbor 104.1.40.2 remote-as 100
neighbor 104.1.40.2 activate
exit-address-family
!
address-family ipv6 vrf l3vpn_vrf11
redistribute connected
neighbor 1411::2 remote-as 100
neighbor 1411::2 activate
neighbor 1412::2 remote-as 100
neighbor 1412::2 activate
neighbor 1413::2 remote-as 100
neighbor 1413::2 activate
neighbor 1414::2 remote-as 100
neighbor 1414::2 activate
neighbor 1415::2 remote-as 100
neighbor 1415::2 activate
neighbor 1416::2 remote-as 100
neighbor 1416::2 activate
neighbor 1417::2 remote-as 100
neighbor 1417::2 activate
neighbor 1418::2 remote-as 100
neighbor 1418::2 activate
neighbor 1419::2 remote-as 100
neighbor 1419::2 activate
neighbor 1420::2 remote-as 100
neighbor 1420::2 activate
neighbor 1421::2 remote-as 100
neighbor 1421::2 activate
neighbor 1422::2 remote-as 100
neighbor 1422::2 activate
neighbor 1423::2 remote-as 100
neighbor 1423::2 activate
neighbor 1424::2 remote-as 100
neighbor 1424::2 activate
neighbor 1425::2 remote-as 100
neighbor 1425::2 activate
neighbor 1426::2 remote-as 100
neighbor 1426::2 activate
neighbor 1427::2 remote-as 100
neighbor 1427::2 activate
neighbor 1428::2 remote-as 100
neighbor 1428::2 activate
neighbor 1429::2 remote-as 100
neighbor 1429::2 activate
neighbor 1430::2 remote-as 100
neighbor 1430::2 activate
neighbor 1431::2 remote-as 100
neighbor 1431::2 activate
neighbor 1432::2 remote-as 100
neighbor 1432::2 activate
neighbor 1433::2 remote-as 100
```

```

neighbor 1433::2 activate
neighbor 1434::2 remote-as 100
neighbor 1434::2 activate
neighbor 1435::2 remote-as 100
neighbor 1435::2 activate
neighbor 1436::2 remote-as 100
neighbor 1436::2 activate
neighbor 1437::2 remote-as 100
neighbor 1437::2 activate
neighbor 1438::2 remote-as 100
neighbor 1438::2 activate
neighbor 1439::2 remote-as 100
neighbor 1439::2 activate
neighbor 1440::2 remote-as 100
neighbor 1440::2 activate
exit-address-family
!
exit

```

Validation

PE1

```
PE1#show ip ospf neighbor
```

```
Total number of full neighbors: 1
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 1.1.1.1 | 1 | Full/Backup | 00:00:36 | 11.1.1.2 | xe11 | 0 |

```
PE1#
```

```
PE1#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 22.22.22.22 | xe11 | Passive | OPERATIONAL | 30 | 01:03:24 |
| | 33.33.33.33 | xe11 | Passive | OPERATIONAL | 30 | 01:03:53 |
| | 1.1.1.1 | xe11 | Active | OPERATIONAL | 30 | 01:04:20 |

```
PE1#
```

```
PE1#show mpls forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP
```

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|-------------|---------|----------|-----------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 1 | 2 | - | - | - | - |
| | - | N/A | 01:05:15 | - | - | - | - |
| | Yes 3 | xe11 | No | 11.1.1.2 | - | - | - |
| L> | 2.2.2.2/32 | 9 | 16 | - | - | - | - |
| | - | N/A | 01:05:15 | - | - | - | - |
| | Yes 24967 | xe11 | No | 11.1.1.2 | - | - | - |
| L> | 12.1.1.0/24 | 10 | 2 | - | - | - | - |
| | - | N/A | 01:04:48 | - | - | - | - |
| | Yes 3 | xe11 | No | 11.1.1.2 | - | - | - |
| L> | 13.1.1.0/24 | 2 | 2 | - | - | - | - |

```

- N/A 01:05:15
1
L> Yes 3 xe11 No 11.1.1.2 - -
14.1.1.0/24 11 18 - - -
- N/A 01:04:48
17
L> Yes 24968 xe11 No 11.1.1.2 - -
15.1.1.0/24 3 4 - - -
- N/A 01:05:15
3
L> Yes 24961 xe11 No 11.1.1.2 - -
16.1.1.0/24 4 6 - - -
- N/A 01:05:15
5
L> Yes 24964 xe11 No 11.1.1.2 - -
17.1.1.0/24 5 8 - - -
- N/A 01:05:15
7
L> Yes 24965 xe11 No 11.1.1.2 - -
18.1.1.0/24 12 20 - - -
- N/A 01:04:48
19
L> Yes 24969 xe11 No 11.1.1.2 - -
19.1.1.0/24 6 10 - - -
- N/A 01:05:15
9
L> Yes 24962 xe11 No 11.1.1.2 - -
20.1.1.0/24 13 22 - - -
- N/A 01:04:44
21
L> Yes 24970 xe11 No 11.1.1.2 - -
21.1.1.0/24 15 26 - - -
- N/A 01:04:14
25
L> Yes 24972 xe11 No 11.1.1.2 - -
22.1.1.0/24 16 28 - - -
- N/A 01:04:14
27
L> Yes 24973 xe11 No 11.1.1.2 - -
22.22.22.22/32 14 24 - - -
- N/A 01:04:42
23
L> Yes 24971 xe11 No 11.1.1.2 - -
33.33.33.33/32 7 12 - - -
- N/A 01:05:15
11
L> Yes 24963 xe11 No 11.1.1.2 - -
44.44.44.44/32 8 14 - - -
- N/A 01:05:15
13
L> Yes 24966 xe11 No 11.1.1.2 - -
55.55.55.55/32 17 30 - - -
- N/A 01:04:14
29
PE1# Yes 24974 xe11 No 11.1.1.2 - -

```

```
PE1#show mpls vpls mesh
```

```
(m) - Service mapped over multipath transport
```

```
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-------------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 11 | 22.22.22.22 | 24971 | 26270 | xe11 | 26240 | 2/Up | 1 L |
| DP | Active | 01:05:05 | - | | | | |
| 11 | 33.33.33.33 | 24963 | 26240 | xe11 | 26240 | 2/Up | 2 L |
| DP | Active | 01:05:54 | - | | | | |

```
PE1#
```


P1

```
P1#show ip ospf neighbor
```

```
Total number of full neighbors: 3
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|----------|-----------|-------------|
| 11.11.11.11 | 1 | Full/DR | 00:00:33 | 11.1.1.1 | xe11 | 0 |
| 22.22.22.22 | 1 | Full/DR | 00:00:39 | 12.1.1.1 | xe7 | 0 |
| 33.33.33.33 | 1 | Full/DR | 00:00:34 | 13.1.1.1 | ce49 | 0 |

```
P1#
```

```
P1#show ldp session
```

```
Codes: m - MD5 password is not set/unset.  
g - GR configuration not set/unset.  
t - TCP MSS not set/unset.  
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 33.33.33.33 | ce49 | Passive | OPERATIONAL | 30 | 01:04:03 |
| | 22.22.22.22 | xe7 | Passive | OPERATIONAL | 30 | 01:03:32 |
| | 11.11.11.11 | xe11 | Passive | OPERATIONAL | 30 | 01:04:18 |

```
P1#
```

```
P1#show mpls forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup  
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,  
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,  
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN  
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP
```

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|----------------|---------|----------|-----------|----------|-----------|------|
| Intf | ELC | NextHop | Algo-Num | UpTime | | | |
| L> | 2.2.2.2/32 | 2 | 5 | - | - | - | - |
| | - | N/A | 01:04:59 | - | - | - | - |
| | Yes | 25601 | ce49 | No | 13.1.1.1 | - | - |
| | | | 44 | - | - | - | - |
| | Yes | 25601 | xe7 | No | 12.1.1.1 | - | - |
| L> | 11.11.11.11/32 | 1 | 2 | - | - | - | - |
| | - | N/A | 01:05:16 | - | - | - | - |
| | Yes | 0 | xe11 | No | 11.1.1.1 | - | - |
| L> | 14.1.1.0/24 | 10 | 27 | - | - | - | - |
| | - | N/A | 01:04:29 | - | - | - | - |
| | Yes | 0 | xe7 | No | 12.1.1.1 | - | - |
| L> | 15.1.1.0/24 | 3 | 8 | - | - | - | - |
| | - | N/A | 01:04:59 | - | - | - | - |
| | Yes | 0 | ce49 | No | 13.1.1.1 | - | - |
| L> | 16.1.1.0/24 | 4 | 11 | - | - | - | - |
| | - | N/A | 01:04:59 | - | - | - | - |
| | Yes | 25604 | ce49 | No | 13.1.1.1 | - | - |
| | | | 45 | - | - | - | - |
| | Yes | 25606 | xe7 | No | 12.1.1.1 | - | - |
| L> | 17.1.1.0/24 | 5 | 14 | - | - | - | - |
| | - | N/A | 01:04:59 | - | - | - | - |
| | Yes | 25605 | ce49 | No | 13.1.1.1 | - | - |
| | | | 46 | - | - | - | - |
| | Yes | 25607 | xe7 | No | 12.1.1.1 | - | - |
| L> | 18.1.1.0/24 | 11 | 27 | - | - | - | - |
| | - | N/A | 01:04:29 | - | - | - | - |
| | Yes | 0 | xe7 | No | 12.1.1.1 | - | - |
| L> | 19.1.1.0/24 | 6 | 8 | - | - | - | - |
| | - | N/A | 01:04:59 | - | - | - | - |

```

      Yes    0          ce49    7    No    -    13.1.1.1    -    -
L> 20.1.1.0/24      9      N/A      8    -    -    -    -
      -          N/A      01:04:44
      Yes    0          ce49    7    No    -    13.1.1.1    -    -
      Yes    0          xe7     26   No    -    12.1.1.1    -    -
L> 21.1.1.0/24     13      N/A      8    -    -    -    -
      -          N/A      01:04:15
      Yes    0          ce49    7    No    -    13.1.1.1    -    -
L> 22.1.1.0/24     14      N/A     27   -    -    -    -
      -          N/A      01:04:15
      Yes    0          xe7     26   No    -    12.1.1.1    -    -
L> 22.22.22.22/32  12      N/A     27   -    -    -    -
      -          N/A      01:04:29
      Yes    0          xe7     26   No    -    12.1.1.1    -    -
L> 33.33.33.33/32  7      N/A      8    -    -    -    -
      -          N/A      01:04:59
      Yes    0          ce49    7    No    -    13.1.1.1    -    -
L> 44.44.44.44/32  8      N/A     17   -    -    -    -
      -          N/A      01:04:59
      Yes    25606      ce49    16   No    -    13.1.1.1    -    -
      Yes    25610      xe7     47   No    -    12.1.1.1    -    -
L> 55.55.55.55/32  15      N/A     42   -    -    -    -
      -          N/A      01:04:15
      Yes    25612      ce49    41   No    -    13.1.1.1    -    -
      Yes    25612      xe7     48   No    -    12.1.1.1    -    -
P1#

```

PE2

```

PE2#show ip ospf neighbor
Total number of full neighbors: 5
OSPF process 100 VRF(default):
Neighbor ID    Pri   State           Dead Time   Address        Interface      Instance ID
1.1.1.1        1     Full/Backup     00:00:30   12.1.1.2      xe12           0
2.2.2.2        1     Full/Backup     00:00:36   14.1.1.2      xe11           0
44.44.44.44    1     Full/DR         00:00:31   18.1.1.2      xe69           0
33.33.33.33    1     Full/DR         00:00:38   20.1.1.2      po10           0
55.55.55.55    1     Full/DR         00:00:37   22.1.1.2      xe70           0
PE2#

```

```

PE2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

Code  Peer IP Address    IF Name    My Role    State        KeepAlive  UpTime
11.11.11.11         xe12       Active     OPERATIONAL 30    01:03:26
33.33.33.33         po10       Passive   OPERATIONAL 30    01:03:24
2.2.2.2             xe11       Active     OPERATIONAL 30    01:03:35
1.1.1.1             xe12       Active     OPERATIONAL 30    01:03:33
44.44.44.44         xe69       Passive   OPERATIONAL 30    01:03:33
55.55.55.55         xe70       Passive   OPERATIONAL 30    01:03:05
PE2#

```

```

PE2#show ip interface bvi11 brief

```

'*' - address is assigned by dhcp client

| Interface | IP-Address | Admin-Status | Link-Status |
|-----------|-------------|--------------|-------------|
| bvi11 | 100.1.11.10 | up | up |

PE2#

Note: BVI interface is up with an ip address on it. If VPLS instance is not attached then the interface will not come up.

PE2#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhife-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|----------------|---------|----------|-----------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 11 | 94 | - | - | - | - |
| | - | N/A | 01:28:56 | 18 | - | - | - |
| | Yes 3 | xe12 | No | 12.1.1.2 | - | - | - |
| L> | 2.2.2.2/32 | 12 | 69 | - | - | - | - |
| | - | N/A | 01:28:56 | 14 | - | - | - |
| | Yes 3 | xe11 | No | 14.1.1.2 | - | - | - |
| L> | 11.1.1.0/24 | 13 | 94 | - | - | - | - |
| | - | N/A | 01:28:56 | 18 | - | - | - |
| | Yes 3 | xe12 | No | 12.1.1.2 | - | - | - |
| L> | 11.11.11.11/32 | 14 | 96 | - | - | - | - |
| | - | N/A | 01:28:56 | 21 | - | - | - |
| | Yes 24960 | xe12 | No | 12.1.1.2 | - | - | - |
| L> | 13.1.1.0/24 | 15 | 94 | - | - | - | - |
| | - | N/A | 01:28:56 | 18 | - | - | - |
| | Yes 3 | xe12 | No | 12.1.1.2 | - | - | - |
| | Yes 0 | po10 | No | 20.1.1.2 | - | - | - |
| L> | 15.1.1.0/24 | 16 | 69 | - | - | - | - |
| | - | N/A | 01:28:56 | 14 | - | - | - |
| | Yes 3 | xe11 | No | 14.1.1.2 | - | - | - |
| | Yes 0 | po10 | No | 20.1.1.2 | - | - | - |
| L> | 16.1.1.0/24 | 17 | 69 | - | - | - | - |
| | - | N/A | 01:28:56 | 14 | - | - | - |
| | Yes 3 | xe11 | No | 14.1.1.2 | - | - | - |
| | Yes 3 | xe69 | No | 18.1.1.2 | - | - | - |
| L> | 17.1.1.0/24 | 18 | 98 | - | - | - | - |
| | - | N/A | 01:28:56 | 24 | - | - | - |
| | Yes 3 | xe69 | No | 18.1.1.2 | - | - | - |
| L> | 19.1.1.0/24 | 19 | 98 | - | - | - | - |
| | - | N/A | 01:28:56 | 24 | - | - | - |
| | Yes 3 | xe69 | No | 18.1.1.2 | - | - | - |
| | Yes 0 | po10 | No | 20.1.1.2 | - | - | - |
| L> | 21.1.1.0/24 | 20 | 61 | - | - | - | - |
| | - | N/A | 01:28:56 | 60 | - | - | - |
| | Yes 3 | xe70 | No | 22.1.1.2 | - | - | - |
| | Yes 0 | po10 | No | 20.1.1.2 | - | - | - |
| L> | 33.33.33.33/32 | 21 | 100 | - | - | - | - |

```

- N/A 01:28:56
27 - 20.1.1.2 - -
L> Yes 0 po10 No 44.44.44.44/32 24 98 - - - -
- N/A 00:00:04 24 -
L> Yes 3 xe69 No 55.55.55.55/32 25 61 - 18.1.1.2 - - -
- N/A 01:28:56 60 -
B> Yes 3 xe70 No 1641::/64 1 84 - 22.1.1.2 - -
No 2.2.2.2 N/A 01:28:52 Yes 24961 -
B> 1642::/64 2 85 - Yes 24962 -
No 2.2.2.2 N/A 01:28:52
B> 1643::/64 3 86 - Yes 24963 -
No 2.2.2.2 N/A 01:28:52
B> 1644::/64 4 87 - Yes 24964 -
No 2.2.2.2 N/A 01:28:52
B> 1645::/64 5 88 - Yes 24965 -
No 2.2.2.2 N/A 01:28:52
B> 1646::/64 6 89 - Yes 24966 -
No 2.2.2.2 N/A 01:28:52
B> 1647::/64 7 90 - Yes 24967 -
No 2.2.2.2 N/A 01:28:52
B> 1648::/64 8 91 - Yes 24968 -
No 2.2.2.2 N/A 01:28:52
B> 1649::/64 9 92 - Yes 24969 -
No 2.2.2.2 N/A 01:28:52
B> 1650::/64 10 93 - Yes 24970 -
No 2.2.2.2 N/A 01:28:52
PE2#

```

PE2#show mpls vrf-forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN

(m) - Service mapped over multipath transport

(e) - Service mapped over ECMP

(D) - Down

Ext-Color - Extended-community color advertised by BGP

B(x) - BGP EVPN MPLS Services

| Code | FEC | FTN-ID | VRF-ID | Nhlfe-ID | Pri | Out-Label | Out- |
|------|---------------|--------|-----------|----------|-----|-----------|------|
| Intf | Nexthop | UpTime | Ext-Color | | | | |
| B> | 104.1.11.0/24 | 1 | 2 | 83 | - | - | - |
| | 01:04:04 | - | | 3 | Yes | 24960 | - |
| | 2.2.2.2 | - | - | | | | |
| B> | 104.1.12.0/24 | 2 | 2 | 83 | - | - | - |
| | 01:04:04 | - | | 3 | Yes | 24960 | - |
| | 2.2.2.2 | - | - | | | | |
| B> | 104.1.13.0/24 | 3 | 2 | 83 | - | - | - |
| | 01:04:04 | - | | 3 | Yes | 24960 | - |
| | 2.2.2.2 | - | - | | | | |
| B> | 104.1.14.0/24 | 4 | 2 | 83 | - | - | - |
| | 01:04:04 | - | | 3 | Yes | 24960 | - |
| | 2.2.2.2 | - | - | | | | |
| B> | 104.1.15.0/24 | 5 | 2 | 83 | - | - | - |
| | 01:04:04 | - | | 3 | Yes | 24960 | - |
| | 2.2.2.2 | - | - | | | | |
| B> | 104.1.16.0/24 | 6 | 2 | 83 | - | - | - |
| | 01:04:04 | - | | 3 | Yes | 24960 | - |
| | 2.2.2.2 | - | - | | | | |
| B> | 104.1.17.0/24 | 7 | 2 | 83 | - | - | - |
| | 01:04:04 | - | | 3 | Yes | 24960 | - |
| | 2.2.2.2 | - | - | | | | |

| | | | | | | | | |
|----|---------------|----|---|----|-----|-------|---|---|
| B> | 104.1.18.0/24 | 8 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.19.0/24 | 9 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.20.0/24 | 10 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.21.0/24 | 11 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.22.0/24 | 12 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.23.0/24 | 13 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.24.0/24 | 14 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.25.0/24 | 15 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.26.0/24 | 16 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.27.0/24 | 17 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.28.0/24 | 18 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.29.0/24 | 19 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.30.0/24 | 20 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.31.0/24 | 21 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.32.0/24 | 22 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.33.0/24 | 23 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.34.0/24 | 53 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | 2.2.2.2 | - | | 3 | Yes | 24960 | - | |
| B> | 104.1.35.0/24 | 54 | 2 | 83 | - | - | - | - |
| | 01:04:04 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |

| | | | | | | | | | |
|----|----------------|----|---|----|-----|-------|---|---|---|
| | 2.2.2.2 | - | - | | | | | | |
| B> | 104.1.136.0/24 | 55 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 104.1.137.0/24 | 56 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 104.1.138.0/24 | 57 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 104.1.139.0/24 | 58 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 104.1.140.0/24 | 59 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 104.1.100.0/24 | 60 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1411::/64 | 24 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1412::/64 | 25 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1413::/64 | 26 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1414::/64 | 27 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1415::/64 | 28 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1416::/64 | 29 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1417::/64 | 30 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1418::/64 | 31 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1419::/64 | 32 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1420::/64 | 33 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1421::/64 | 34 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |
| | | | | 3 | Yes | 24960 | - | | |
| | 2.2.2.2 | - | - | | | | | | |
| B> | 1422::/64 | 35 | 2 | 83 | - | - | - | - | - |
| | 01:04:04 | - | | | | | | | |

| | | | | | | | | | |
|----|-----------|----------|---|----|---|-----|-------|-------|---|
| | | 2.2.2.2 | - | - | 3 | Yes | 24960 | - | |
| B> | 1423::/64 | 01:04:04 | - | 36 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1424::/64 | 01:04:04 | - | 37 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1425::/64 | 01:04:04 | - | 38 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1426::/64 | 01:04:04 | - | 39 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1427::/64 | 01:04:04 | - | 40 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1428::/64 | 01:04:04 | - | 41 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1429::/64 | 01:04:04 | - | 42 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1430::/64 | 01:04:04 | - | 43 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1431::/64 | 01:04:04 | - | 44 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1432::/64 | 01:04:04 | - | 45 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1433::/64 | 01:04:04 | - | 46 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1434::/64 | 01:04:04 | - | 47 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1435::/64 | 01:04:04 | - | 48 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1436::/64 | 01:04:04 | - | 49 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1437::/64 | 01:04:04 | - | 50 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1438::/64 | 01:04:04 | - | 51 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1439::/64 | 01:04:04 | - | 52 | 2 | 83 | - | - | - |
| | | | | | | 3 | Yes | 24960 | - |
| B> | 1440::/64 | | - | 61 | 2 | 83 | - | - | - |

```

01:04:04 -
2.2.2.2 - - 3 Yes 24960 -
PE2#

```

```

PE2#show ipv6 bgp summary
BGP router identifier 22.22.22.22, local AS number 123456
BGP table version is 12
2 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|-----|--------|--------|--------|-----|------|----------|--------------|------|
| 641::11 | 4 | 100 | 4 | 5 | 10 | 0 | 0 | 00:00:13 | 10 | |
| 642::11 | 4 | 100 | 4 | 6 | 12 | 0 | 0 | 00:00:03 | 10 | |
| 643::11 | 4 | 100 | 5 | 8 | 9 | 0 | 0 | 00:00:46 | 10 | |
| 644::11 | 4 | 100 | 4 | 6 | 11 | 0 | 0 | 00:00:10 | 10 | |
| 645::11 | 4 | 100 | 5 | 9 | 9 | 0 | 0 | 00:00:42 | 10 | |
| 646::11 | 4 | 100 | 5 | 9 | 9 | 0 | 0 | 00:00:42 | 10 | |
| 647::11 | 4 | 100 | 5 | 9 | 11 | 0 | 0 | 00:00:38 | 10 | |
| 648::11 | 4 | 100 | 5 | 9 | 10 | 0 | 0 | 00:00:38 | 10 | |
| 649::11 | 4 | 100 | 5 | 9 | 11 | 0 | 0 | 00:00:34 | 10 | |
| 650::11 | 4 | 100 | 5 | 8 | 9 | 0 | 0 | 00:00:43 | 10 | |

Total number of neighbors 10

Total number of Established sessions 10

PE2#

Note: The above ipv6 neighbors are 6pe neighbors.

```

PE2#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | |
|--------------|-------------|--------------|-----------|--------------|-----------|---------|----------|---|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | | |
| 11 | 11.11.11.11 | 24960 | 26240 | xe12 | 26270 | 2/Up | 1 | L |
| DP | Active | 01:05:06 | - | | | | | |
| 11 | 33.33.33.33 | 0 | 26270 | po10 | 26280 | 2/Up | 2 | L |
| DP | Active | 01:05:06 | - | | | | | |

PE2#

```

PE2#show mpls vpls ldp-vpls11
Virtual Private LAN Service Instance: ldp-vpls11, ID: 11
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Enabled
Flow Label Status: Enabled, Direction: Both, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

```

Configured interfaces:

```

Interface: bv111
Status: Up
Subinterface Match Criteria(s) :
dot1q 11

```

```

Interface: xe73.11
Status: Up
Subinterface Match Criteria(s) :
dot1q 11

```

Mesh Peers:

```

11.11.11.11 (Type: Ethernet VLAN) (Negotiated - CW: Yes, FAT: Both) (Up) (UpTime: 01:05:15)
FEC signaling element: FEC128

```



```

33.33.33.33 (Type: Ethernet) (Negotiated - CW: Yes, FAT: Both) (Up) (UpTime: 01:05:15)
FEC signaling element: FEC128
Spoke Peers:
vc11 (Dn) (Reason: VC on standby)

```

PE2#

Note: Here we can see that BVI interface is attached to VPLS session. If VPLS instance is not attached to a BVI, then the interface will not come up.

This shows BVI is able bind with VPLS(L2vpn) and also VRF(L3vpn) domains to stitch the traffic.

```

PE2#show vrrp summary
VRRP Version: 3
VMAC enabled
Backward Compatibility disabled

```

| Interface | Admin | IpVersion | VRGroup | VRState | Priority | VIP | InterfaceIp |
|-----------|-------|-----------|---------|---------|----------|----------------|---------------------------|
| bvi11 | Up | IPv4 | 1 | Master | 200 | 100.1.11.100 | 100.1.11.10 |
| bvi11 | Up | IPv6 | 1 | Master | 200 | fe80::11:0:0:1 | fe80::e69d:73ff:feb1:c301 |

Total Number of IPv4 VRRP group 1
Total Number of IPv6 VRRP group 1
PE2#

Note: Here we will have VRRP Virtual-ip as the gateway for the traffic coming from L2VPN domain while the destination is on L3VPN domain. Both vrrp-v4 and vrrp-v6 is configured as shown.

PE2 is the Master Vrrp with highest priority.

```

PE2#show ip bgp vpnv4 all summary
BGP router identifier 22.22.22.22, local AS number 123456
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|--------|--------|--------|--------|-----|------|----------|--------------|------|
| 2.2.2.2 | 4 | 123456 | 178 | 164 | 5 | 0 | 0 | 01:05:56 | 61 | |

Total number of neighbors 1

Total number of Established sessions 1

PE3

```
PE3#show ip ospf neighbor
```

```

Total number of full neighbors: 5
OSPF process 100 VRF(default):

```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 1.1.1.1 | 1 | Full/Backup | 00:00:38 | 13.1.1.2 | ce6 | 0 |
| 2.2.2.2 | 1 | Full/Backup | 00:00:34 | 15.1.1.2 | ce9/3 | 0 |
| 44.44.44.44 | 1 | Full/DR | 00:00:37 | 19.1.1.2 | ce9/4 | 0 |
| 22.22.22.22 | 1 | Full/Backup | 00:00:29 | 20.1.1.1 | po10 | 0 |
| 55.55.55.55 | 1 | Full/DR | 00:00:35 | 21.1.1.2 | xe3 | 0 |

PE3#

```
PE3#show ldp session
```

```

Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 11.11.11.11 | ce6 | Active | OPERATIONAL | 30 | 01:03:53 |
| | 22.22.22.22 | po10 | Active | OPERATIONAL | 30 | 01:03:24 |
| | 1.1.1.1 | ce6 | Active | OPERATIONAL | 30 | 01:04:03 |
| | 2.2.2.2 | ce9/3 | Active | OPERATIONAL | 30 | 01:04:03 |
| | 44.44.44.44 | ce9/4 | Passive | OPERATIONAL | 30 | 01:04:03 |

```

55.55.55.55      xe3      Passive  OPERATIONAL  30      01:03:37
PE3#

```

```
PE3#show mpls forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|----------------|---------|----------|-----------|-----|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 11 | 15 | - | - | - | - |
| | - | N/A | 01:05:00 | - | - | - | - |
| | Yes 3 | ce6 | 14 | No | - | - | - |
| L> | 2.2.2.2/32 | 12 | 43 | - | - | - | - |
| | - | N/A | 01:05:00 | - | - | - | - |
| | Yes 3 | ce9/3 | 17 | No | - | - | - |
| L> | 11.1.1.0/24 | 13 | 15 | - | - | - | - |
| | - | N/A | 01:05:00 | - | - | - | - |
| | Yes 3 | ce6 | 14 | No | - | - | - |
| L> | 11.11.11.11/32 | 14 | 22 | - | - | - | - |
| | - | N/A | 01:05:00 | - | - | - | - |
| | Yes 24960 | ce6 | 21 | No | - | - | - |
| L> | 12.1.1.0/24 | 19 | 15 | - | - | - | - |
| | - | N/A | 01:04:48 | - | - | - | - |
| | Yes 3 | ce6 | 14 | No | - | - | - |
| | Yes 0 | po10 | 45 | No | - | - | - |
| L> | 14.1.1.0/24 | 20 | 43 | - | - | - | - |
| | - | N/A | 01:04:48 | - | - | - | - |
| | Yes 3 | ce9/3 | 17 | No | - | - | - |
| | Yes 0 | po10 | 45 | No | - | - | - |
| L> | 16.1.1.0/24 | 15 | 25 | - | - | - | - |
| | - | N/A | 01:05:00 | - | - | - | - |
| | Yes 3 | ce9/4 | 24 | No | - | - | - |
| | Yes 3 | ce9/3 | 17 | No | - | - | - |
| L> | 17.1.1.0/24 | 16 | 25 | - | - | - | - |
| | - | N/A | 01:05:00 | - | - | - | - |
| | Yes 3 | ce9/4 | 24 | No | - | - | - |
| L> | 18.1.1.0/24 | 21 | 25 | - | - | - | - |
| | - | N/A | 01:04:48 | - | - | - | - |
| | Yes 3 | ce9/4 | 24 | No | - | - | - |
| | Yes 0 | po10 | 45 | No | - | - | - |
| L> | 22.1.1.0/24 | 23 | 40 | - | - | - | - |
| | - | N/A | 01:04:34 | - | - | - | - |
| | Yes 3 | xe3 | 39 | No | - | - | - |
| | Yes 0 | po10 | 45 | No | - | - | - |
| L> | 22.22.22.22/32 | 26 | 52 | - | - | - | - |
| | - | N/A | 01:04:21 | - | - | - | - |
| | Yes 0 | po10 | 45 | No | - | - | - |
| L> | 44.44.44.44/32 | 17 | 25 | - | - | - | - |
| | - | N/A | 00:09:57 | - | - | - | - |
| L> | 55.55.55.55/32 | 24 | 40 | - | - | - | - |
| | - | N/A | 01:04:34 | - | - | - | - |

```

          Yes    3          xe3          39          -          21.1.1.2          -          -
B> 650::/64      25          54          -          Yes    24978          -
    No          2.2.2.2          N/A          01:03:54
B> 1641::/64     1          4          -          Yes    24961          -
    No          2.2.2.2          N/A          01:05:00
B> 1642::/64     2          5          -          Yes    24962          -
    No          2.2.2.2          N/A          01:05:00
B> 1643::/64     3          6          -          Yes    24963          -
    No          2.2.2.2          N/A          01:05:00
B> 1644::/64     4          7          -          Yes    24964          -
    No          2.2.2.2          N/A          01:05:00
B> 1645::/64     5          8          -          Yes    24965          -
    No          2.2.2.2          N/A          01:05:00
B> 1646::/64     6          9          -          Yes    24966          -
    No          2.2.2.2          N/A          01:05:00
B> 1647::/64     7          10         -          Yes    24967          -
    No          2.2.2.2          N/A          01:05:00
B> 1648::/64     8          11         -          Yes    24968          -
    No          2.2.2.2          N/A          01:05:00
B> 1649::/64     9          12         -          Yes    24969          -
    No          2.2.2.2          N/A          01:05:00
B> 1650::/64    10          13         -          Yes    24970          -
    No          2.2.2.2          N/A          01:05:00
PE3#

```

PE3#show mpls vrf-forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN

(m) - Service mapped over multipath transport

(e) - Service mapped over ECMP

(D) - Down

Ext-Color - Extended-community color advertised by BGP

B(x) - BGP EVPN MPLS Services

| Code | FEC | Nexthop | FTN-ID | VRF-ID | Nhlfe-ID | Pri | Out-Label | Out- |
|------|---------------|---------|--------|--------|-----------|-------|-----------|------|
| Intf | | | UpTime | | Ext-Color | | | |
| B> | 104.1.11.0/24 | 1 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | | | | | |
| B> | 104.1.12.0/24 | 2 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | | | | | |
| B> | 104.1.13.0/24 | 3 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | | | | | |
| B> | 104.1.14.0/24 | 4 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | | | | | |
| B> | 104.1.15.0/24 | 5 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | | | | | |
| B> | 104.1.16.0/24 | 6 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | | | | | |
| B> | 104.1.17.0/24 | 7 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | | | | | |
| B> | 104.1.18.0/24 | 8 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | | | | | |
| B> | 104.1.19.0/24 | 9 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | | | | | |

| | | | | | | | | |
|----|---------------|----|---|----|-----|-------|---|---|
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.20.0/24 | 10 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.21.0/24 | 11 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.22.0/24 | 12 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.23.0/24 | 13 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.24.0/24 | 14 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.25.0/24 | 15 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.26.0/24 | 16 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.27.0/24 | 17 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.28.0/24 | 18 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.29.0/24 | 19 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.30.0/24 | 20 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.31.0/24 | 21 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.32.0/24 | 22 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.33.0/24 | 23 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.34.0/24 | 53 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.35.0/24 | 54 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.36.0/24 | 55 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |
| | | | | 3 | Yes | 24960 | - | |
| | 2.2.2.2 | - | - | 19 | - | - | - | - |
| B> | 104.1.37.0/24 | 56 | 2 | 19 | - | - | - | - |
| | 01:05:06 | - | | | | | | |

| | | | | | | | | | |
|----|----------------|----------|----|---|----|-----|-------|---|---|
| | | 2.2.2.2 | - | - | 3 | Yes | 24960 | - | |
| B> | 104.1.38.0/24 | 01:05:06 | 57 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 104.1.39.0/24 | 01:05:06 | 58 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 104.1.40.0/24 | 01:05:06 | 59 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 104.1.100.0/24 | 01:05:06 | 60 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1411::/64 | 01:05:06 | 24 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1412::/64 | 01:05:06 | 25 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1413::/64 | 01:05:06 | 26 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1414::/64 | 01:05:06 | 27 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1415::/64 | 01:05:06 | 28 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1416::/64 | 01:05:06 | 29 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1417::/64 | 01:05:06 | 30 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1418::/64 | 01:05:06 | 31 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1419::/64 | 01:05:06 | 32 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1420::/64 | 01:05:06 | 33 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1421::/64 | 01:05:06 | 34 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1422::/64 | 01:05:06 | 35 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1423::/64 | 01:05:06 | 36 | 2 | 19 | - | - | - | - |
| | | | | | 3 | Yes | 24960 | - | |
| B> | 1424::/64 | | 37 | 2 | 19 | - | - | - | - |

```

01:05:06 -
      2.2.2.2 - 3 Yes 24960 -
B> 1425::/64 38 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 39 2 19 - - -
B> 1426::/64 39 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 40 2 19 - - -
B> 1427::/64 40 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 41 2 19 - - -
B> 1428::/64 41 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 42 2 19 - - -
B> 1429::/64 42 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 43 2 19 - - -
B> 1430::/64 43 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 44 2 19 - - -
B> 1431::/64 44 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 45 2 19 - - -
B> 1432::/64 45 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 46 2 19 - - -
B> 1433::/64 46 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 47 2 19 - - -
B> 1434::/64 47 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 48 2 19 - - -
B> 1435::/64 48 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 49 2 19 - - -
B> 1436::/64 49 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 50 2 19 - - -
B> 1437::/64 50 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 51 2 19 - - -
B> 1438::/64 51 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 52 2 19 - - -
B> 1439::/64 52 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - 61 2 19 - - -
B> 1440::/64 61 2 19 - - -
    01:05:06 - 3 Yes 24960 -
      2.2.2.2 - - -
PE3#

```

```
PE3#show ip interface bvi11 brief
```

'*' - address is assigned by dhcp client

| Interface | IP-Address | Admin-Status | Link-Status |
|-----------|-------------|--------------|-------------|
| bvi11 | 100.1.11.20 | up | up |

PE3#

PE3#show mpls vpls mesh
 (m) - Service mapped over multipath transport
 (e) - Service mapped over LDP ECMP

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-------------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 11 | 11.11.11.11 | 24960 | 26240 | ce6 | 26240 | 2/Up | 1 L |
| DP | Active | 01:05:53 | - | | | | |
| 11 | 22.22.22.22 | 0 | 26280 | po10 | 26270 | 2/Up | 2 L |
| DP | Active | 01:05:05 | - | | | | |

PE3#

PE3#show mpls vpls ldp-vpls11
 Virtual Private LAN Service Instance: ldp-vpls11, ID: 11
 SIG-Protocol: LDP
 Attachment-Circuit: UP
 Learning: Enabled
 Control-Word: Enabled
 Flow Label Status: Enabled, Direction: Both, Static: No
 Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
 Description: none
 service-tpid: dot1q
 Operating mode: Raw
 MAC Withdrawal:

Configured interfaces:
 Interface: bvi11
 Status: Up
 Subinterface Match Criteria(s) :
 dot1q 11

Interface: ce9/2.11
 Status: Up
 Subinterface Match Criteria(s) :
 dot1q 11

Mesh Peers:
 11.11.11.11 (Type: Ethernet) (Negotiated - CW: Yes, FAT: Both) (Up) (UpTime: 01:06:11)
 FEC signaling element: FEC128
 22.22.22.22 (Type: Ethernet) (Negotiated - CW: Yes, FAT: Both) (Up) (UpTime: 01:05:23)
 FEC signaling element: FEC128
 Spoke Peers:
 vc111 (Up) (UpTime 01:05:55)

PE3#

PE3#show vrrp summary
 VRRP Version: 3
 VMAC enabled
 Backward Compatibility disabled

| Interface | Admin | IpVersion | VRGroup | VRState | Priority | VIP | InterfaceIp |
|-----------|-------|-----------|---------|---------|----------|----------------|---------------------------|
| bvi11 | Up | IPv4 | 1 | Backup | 150 | 100.1.11.100 | 100.1.11.20 |
| bvi11 | Up | IPv6 | 1 | Backup | 150 | fe80::11:0:0:1 | fe80::5e07:58ff:fe51:cf01 |

Total Number of IPv4 VRRP group 30
 Total Number of IPv6 VRRP group 40
 PE3#

PE3#show ip bgp vpnv4 all summary
 BGP router identifier 33.33.33.33, local AS number 123456

```
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|--------|--------|--------|--------|-----|------|----------|--------------|------|
| 2.2.2.2 | 4 | 123456 | 176 | 162 | 5 | 0 | 0 | 01:07:10 | | 61 |

Total number of neighbors 1

Total number of Established sessions 1
PE3#

P2

```
RR-7034#show ip ospf neighbor
```

```
Total number of full neighbors: 3
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|----------|-----------|-------------|
| 22.22.22.22 | 1 | Full/DR | 00:00:33 | 14.1.1.1 | xe3 | 0 |
| 33.33.33.33 | 1 | Full/DR | 00:00:32 | 15.1.1.1 | xe48 | 0 |
| 44.44.44.44 | 1 | Full/DR | 00:00:34 | 16.1.1.1 | xe8 | 0 |

```
RR-7034#
```

```
RR-7034#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 22.22.22.22 | xe3 | Passive | OPERATIONAL | 30 | 01:03:34 |
| | 44.44.44.44 | xe8 | Passive | OPERATIONAL | 30 | 01:04:24 |
| | 33.33.33.33 | xe48 | Passive | OPERATIONAL | 30 | 01:04:03 |

```
RR-7034#show mpls forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP
```

```
FTN-ECMP LDP: Disabled, SR: Disabled
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|----------------|---------|----------|-----------|----------|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 24 | 10 | - | - | - | - |
| | - | N/A | 01:04:59 | - | - | - | - |
| | Yes | 25600 | xe48 | No | 15.1.1.1 | - | - |
| | | | 43 | - | - | - | - |
| | Yes | 25600 | xe3 | No | 14.1.1.1 | - | - |
| L> | 11.1.1.0/24 | 25 | 13 | - | - | - | - |
| | - | N/A | 01:04:59 | - | - | - | - |
| | Yes | 25602 | xe48 | No | 15.1.1.1 | - | - |
| | | | 44 | - | - | - | - |
| | Yes | 25602 | xe3 | No | 14.1.1.1 | - | - |
| L> | 11.11.11.11/32 | 26 | 16 | - | - | - | - |
| | - | N/A | 01:04:59 | - | - | - | - |
| | Yes | 25603 | xe48 | No | 15.1.1.1 | - | - |
| | | | 45 | - | - | - | - |
| | Yes | 25603 | xe3 | No | 14.1.1.1 | - | - |
| L> | 12.1.1.0/24 | 35 | 37 | - | - | - | - |
| | - | N/A | 01:04:34 | - | - | - | - |
| | | | 36 | - | - | - | - |
| | Yes | 0 | xe3 | No | 14.1.1.1 | - | - |
| L> | 13.1.1.0/24 | 27 | 19 | - | - | - | - |
| | - | N/A | 01:04:59 | - | - | - | - |

| | | | | | | | | | | |
|----|----------------|-------------|------|----------|----|----------|----------|-------|---|---|
| | Yes | 0 | xe48 | 18 | No | - | 15.1.1.1 | - | - | - |
| L> | 17.1.1.0/24 | 12 | | 3 | | - | - | - | - | - |
| | - | N/A | | 01:05:20 | | - | | | | |
| | Yes | 3 | xe8 | 2 | No | - | 16.1.1.1 | - | - | - |
| L> | 18.1.1.0/24 | 29 | | 3 | | - | - | - | - | - |
| | - | N/A | | 01:04:48 | | - | | | | |
| | Yes | 3 | xe8 | 2 | No | - | 16.1.1.1 | - | - | - |
| | Yes | 0 | xe3 | 36 | No | - | 14.1.1.1 | - | - | - |
| L> | 19.1.1.0/24 | 11 | | 3 | | - | - | - | - | - |
| | - | N/A | | 01:05:20 | | - | | | | |
| | Yes | 3 | xe8 | 2 | No | - | 16.1.1.1 | - | - | - |
| | Yes | 0 | xe48 | 18 | No | - | 15.1.1.1 | - | - | - |
| L> | 20.1.1.0/24 | 30 | | 19 | | - | - | - | - | - |
| | - | N/A | | 01:04:44 | | - | | | | |
| | Yes | 0 | xe48 | 18 | No | - | 15.1.1.1 | - | - | - |
| | Yes | 0 | xe3 | 36 | No | - | 14.1.1.1 | - | - | - |
| L> | 21.1.1.0/24 | 31 | | 19 | | - | - | - | - | - |
| | - | N/A | | 01:04:36 | | - | | | | |
| | Yes | 0 | xe48 | 18 | No | - | 15.1.1.1 | - | - | - |
| L> | 22.1.1.0/24 | 32 | | 37 | | - | - | - | - | - |
| | - | N/A | | 01:04:36 | | - | | | | |
| | Yes | 0 | xe3 | 36 | No | - | 14.1.1.1 | - | - | - |
| L> | 22.22.22.22/32 | 36 | | 37 | | - | - | - | - | - |
| | - | N/A | | 01:04:30 | | - | | | | |
| | Yes | 0 | xe3 | 36 | No | - | 14.1.1.1 | - | - | - |
| L> | 33.33.33.33/32 | 28 | | 19 | | - | - | - | - | - |
| | - | N/A | | 01:04:59 | | - | | | | |
| | Yes | 0 | xe48 | 18 | No | - | 15.1.1.1 | - | - | - |
| L> | 44.44.44.44/32 | 13 | | 3 | | - | - | - | - | - |
| | - | N/A | | 01:05:20 | | - | | | | |
| | Yes | 3 | xe8 | 2 | No | - | 16.1.1.1 | - | - | - |
| L> | 55.55.55.55/32 | 33 | | 30 | | - | - | - | - | - |
| | - | N/A | | 01:04:36 | | - | | | | |
| | Yes | 25612 | xe48 | 29 | No | - | 15.1.1.1 | - | - | - |
| | Yes | 25612 | xe3 | 41 | No | - | 14.1.1.1 | - | - | - |
| B> | 641::/64 | 14 | | 46 | | - | Yes | 28163 | - | - |
| | No | 22.22.22.22 | | N/A | | 01:04:30 | | | | |
| B> | 642::/64 | 15 | | 46 | | - | Yes | 28163 | - | - |
| | No | 22.22.22.22 | | N/A | | 01:04:30 | | | | |
| B> | 643::/64 | 16 | | 46 | | - | Yes | 28163 | - | - |
| | No | 22.22.22.22 | | N/A | | 01:04:30 | | | | |
| B> | 644::/64 | 17 | | 46 | | - | Yes | 28163 | - | - |
| | No | 22.22.22.22 | | N/A | | 01:04:30 | | | | |
| B> | 645::/64 | 18 | | 46 | | - | Yes | 28163 | - | - |
| | No | 22.22.22.22 | | N/A | | 01:04:30 | | | | |
| B> | 646::/64 | 19 | | 46 | | - | Yes | 28163 | - | - |
| | No | 22.22.22.22 | | N/A | | 01:04:30 | | | | |
| B> | 647::/64 | 20 | | 46 | | - | Yes | 28163 | - | - |
| | No | 22.22.22.22 | | N/A | | 01:04:30 | | | | |
| B> | 648::/64 | 21 | | 46 | | - | Yes | 28163 | - | - |
| | No | 22.22.22.22 | | N/A | | 01:04:30 | | | | |
| B> | 649::/64 | 22 | | 46 | | - | Yes | 28163 | - | - |
| | No | 22.22.22.22 | | N/A | | 01:04:30 | | | | |
| B> | 64a::/64 | 23 | | 6 | | - | Yes | 26881 | - | - |
| | No | 33.33.33.33 | | N/A | | 01:04:59 | | | | |
| B> | 650::/64 | 34 | | 46 | | - | Yes | 28163 | - | - |

```

      No      22.22.22.22      N/A      01:03:55
B> 1641::/64      1      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
B> 1642::/64      2      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
B> 1643::/64      3      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
B> 1644::/64      4      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
B> 1645::/64      5      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
B> 1646::/64      6      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
B> 1647::/64      7      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
B> 1648::/64      8      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
B> 1649::/64      9      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
B> 1650::/64     10      1      -      Yes      24961      -
      No      44.44.44.44      N/A      01:05:20
RR-7034#

```

```

RR-7034#show ip bgp vpnv4 all summary
BGP router identifier 2.2.2.2, local AS number 123456
BGP table version is 6
1 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|-------------|---|--------|--------|--------|--------|-----|------|----------|--------------|------|
| 22.22.22.22 | 4 | 123456 | 163 | 178 | 6 | 0 | 0 | 01:05:57 | | 30 |
| 33.33.33.33 | 4 | 123456 | 162 | 177 | 6 | 0 | 0 | 01:07:11 | | 30 |
| 44.44.44.44 | 4 | 123456 | 160 | 179 | 6 | 0 | 0 | 01:07:26 | | 31 |

Total number of neighbors 3

Total number of Established sessions 3

RR-7034#

PE4

```
PE4#show ip ospf neighbor
```

Total number of full neighbors: 3

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 2.2.2.2 | 1 | Full/Backup | 00:00:35 | 16.1.1.2 | xe6 | 0 |
| 22.22.22.22 | 1 | Full/Backup | 00:00:34 | 18.1.1.1 | xe27 | 0 |
| 33.33.33.33 | 1 | Full/Backup | 00:00:31 | 19.1.1.1 | xe26 | 0 |

PE4#

```
PE4#show mpls forwarding-table
```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-ID | Pri | Out-Label | Out- |
|------|------------|---------|----------|-----------|----------|-----------|------|
| Intf | ELC | Nexthop | Algo-Num | UpTime | | | |
| L> | 1.1.1.1/32 | 15 | 21 | - | - | - | - |
| | - | N/A | 01:05:01 | - | | | |
| | | | 20 | - | | | |
| | Yes | 25600 | xe26 | No | 19.1.1.1 | - | - |

| | | | | | | | | | | | |
|----|----------------|---------|-------|------|----------|----|---|----------|-------|---|---|
| | | Yes | 25600 | xe27 | 74 | No | - | 18.1.1.1 | - | - | - |
| L> | 2.2.2.2/32 | | 1 | | 4 | | - | - | - | - | - |
| | - | | N/A | | 01:05:21 | | - | | | | |
| | | Yes | 3 | xe6 | 3 | No | - | 16.1.1.2 | - | - | - |
| L> | 11.1.1.0/24 | | 16 | | 24 | | - | - | - | - | - |
| | - | | N/A | | 01:05:01 | | - | | | | |
| | | Yes | 25602 | xe26 | 23 | No | - | 19.1.1.1 | - | - | - |
| | | Yes | 25602 | xe27 | 75 | No | - | 18.1.1.1 | - | - | - |
| L> | 11.11.11.11/32 | | 17 | | 27 | | - | - | - | - | - |
| | - | | N/A | | 01:05:01 | | - | | | | |
| | | Yes | 25603 | xe26 | 26 | No | - | 19.1.1.1 | - | - | - |
| | | Yes | 25603 | xe27 | 76 | No | - | 18.1.1.1 | - | - | - |
| L> | 12.1.1.0/24 | | 25 | | 77 | | - | - | - | - | - |
| | - | | N/A | | 01:04:35 | | - | | | | |
| | | Yes | 0 | xe27 | 57 | No | - | 18.1.1.1 | - | - | - |
| L> | 13.1.1.0/24 | | 18 | | 30 | | - | - | - | - | - |
| | - | | N/A | | 01:05:01 | | - | | | | |
| | | Yes | 0 | xe26 | 29 | No | - | 19.1.1.1 | - | - | - |
| L> | 14.1.1.0/24 | | 20 | | 4 | | - | - | - | - | - |
| | - | | N/A | | 01:04:49 | | - | | | | |
| | | Yes | 3 | xe6 | 3 | No | - | 16.1.1.2 | - | - | - |
| | | Yes | 0 | xe27 | 57 | No | - | 18.1.1.1 | - | - | - |
| L> | 15.1.1.0/24 | | 2 | | 4 | | - | - | - | - | - |
| | - | | N/A | | 01:05:21 | | - | | | | |
| | | Yes | 3 | xe6 | 3 | No | - | 16.1.1.2 | - | - | - |
| | | Yes | 0 | xe26 | 29 | No | - | 19.1.1.1 | - | - | - |
| L> | 20.1.1.0/24 | | 21 | | 30 | | - | - | - | - | - |
| | - | | N/A | | 01:04:46 | | - | | | | |
| | | Yes | 0 | xe26 | 29 | No | - | 19.1.1.1 | - | - | - |
| | | Yes | 0 | xe27 | 57 | No | - | 18.1.1.1 | - | - | - |
| L> | 21.1.1.0/24 | | 22 | | 30 | | - | - | - | - | - |
| | - | | N/A | | 01:04:36 | | - | | | | |
| | | Yes | 0 | xe26 | 29 | No | - | 19.1.1.1 | - | - | - |
| L> | 22.1.1.0/24 | | 23 | | 77 | | - | - | - | - | - |
| | - | | N/A | | 01:04:36 | | - | | | | |
| | | Yes | 0 | xe27 | 57 | No | - | 18.1.1.1 | - | - | - |
| L> | 22.22.22.22/32 | | 28 | | 77 | | - | - | - | - | - |
| | - | | N/A | | 00:17:23 | | - | | | | |
| | | Yes | 0 | xe27 | 57 | No | - | 18.1.1.1 | - | - | - |
| L> | 33.33.33.33/32 | | 19 | | 30 | | - | - | - | - | - |
| | - | | N/A | | 00:16:54 | | - | | | | |
| | | Yes | 0 | xe26 | 29 | No | - | 19.1.1.1 | - | - | - |
| L> | 55.55.55.55/32 | | 24 | | 41 | | - | - | - | - | - |
| | - | | N/A | | 01:04:36 | | - | | | | |
| | | Yes | 25612 | xe26 | 40 | No | - | 19.1.1.1 | - | - | - |
| | | Yes | 25612 | xe27 | 79 | No | - | 18.1.1.1 | - | - | - |
| B> | 641::/64 | | 3 | | 82 | | - | Yes | 24980 | - | - |
| | No | 2.2.2.2 | | N/A | | | - | 01:05:07 | | | |
| B> | 642::/64 | | 4 | | 35 | | - | Yes | 24982 | - | - |
| | No | 2.2.2.2 | | N/A | | | - | 01:05:07 | | | |
| B> | 643::/64 | | 5 | | 83 | | - | Yes | 24983 | - | - |

```

      No      2.2.2.2      N/A      01:05:07
B> 644::/64      6      84      -      Yes      24972      -
      No      2.2.2.2      N/A      01:05:07
B> 645::/64      7      85      -      Yes      24973      -
      No      2.2.2.2      N/A      01:05:07
B> 646::/64      8      86      -      Yes      24974      -
      No      2.2.2.2      N/A      01:05:07
B> 647::/64      9      87      -      Yes      24975      -
      No      2.2.2.2      N/A      01:05:07
B> 648::/64      10     88      -      Yes      24976      -
      No      2.2.2.2      N/A      01:05:07
B> 649::/64      11     89      -      Yes      24977      -
      No      2.2.2.2      N/A      01:05:07
B> 64a::/64      12     16      -      Yes      24981      -
      No      2.2.2.2      N/A      01:05:07
B> 650::/64      26     90      -      Yes      24978      -
      No      2.2.2.2      N/A      01:03:52
PE4#

```

```

PE4#show ldp session
Codes: m - MD5 password is not set/unset.
      g - GR configuration not set/unset.
      t - TCP MSS not set/unset.
      Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | xe6 | Active | OPERATIONAL | 30 | 01:04:24 |
| | 33.33.33.33 | xe26 | Active | OPERATIONAL | 30 | 01:04:05 |
| | 22.22.22.22 | xe27 | Active | OPERATIONAL | 30 | 01:03:34 |

```
PE4#
```

```

PE4#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

```

| Code | FEC | FTN-ID | VRF-ID | Nhlfe-ID | Pri | Out-Label | Out- |
|------|---------------|--------|-----------|----------|-----|-----------|------|
| Intf | Nexthop | UpTime | Ext-Color | | | | |
| B> | 100.1.11.0/24 | 1 | 2 | 81 | - | - | - |
| | 01:05:15 | - | - | 80 | Yes | 24979 | - |
| B> | 100.1.12.0/24 | 2 | 2 | 81 | - | - | - |
| | 01:05:15 | - | - | 80 | Yes | 24979 | - |
| B> | 100.1.13.0/24 | 3 | 2 | 81 | - | - | - |
| | 01:05:15 | - | - | 80 | Yes | 24979 | - |
| B> | 100.1.14.0/24 | 4 | 2 | 81 | - | - | - |
| | 01:05:15 | - | - | 80 | Yes | 24979 | - |
| B> | 100.1.15.0/24 | 5 | 2 | 81 | - | - | - |
| | 01:05:15 | - | - | 80 | Yes | 24979 | - |
| B> | 100.1.16.0/24 | 6 | 2 | 81 | - | - | - |
| | 01:05:15 | - | - | 80 | Yes | 24979 | - |
| B> | 100.1.17.0/24 | 7 | 2 | 81 | - | - | - |
| | 01:05:15 | - | - | 80 | Yes | 24979 | - |
| B> | 100.1.18.0/24 | 8 | 2 | 81 | - | - | - |
| | 01:05:15 | - | - | | | | |

| | | | | | | | | | |
|----|---------------|----------|---|----|----|-----|-------|-------|---|
| | | 2.2.2.2 | - | - | 80 | Yes | 24979 | - | |
| B> | 100.1.19.0/24 | 01:05:15 | - | 9 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.20.0/24 | 01:05:15 | - | 10 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.21.0/24 | 01:05:15 | - | 11 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.22.0/24 | 01:05:15 | - | 12 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.23.0/24 | 01:05:15 | - | 13 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.24.0/24 | 01:05:15 | - | 14 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.25.0/24 | 01:05:15 | - | 15 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.26.0/24 | 01:05:15 | - | 16 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.27.0/24 | 01:05:15 | - | 17 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.28.0/24 | 01:05:15 | - | 18 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.29.0/24 | 01:05:15 | - | 19 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.30.0/24 | 01:05:15 | - | 20 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.31.0/24 | 01:05:15 | - | 21 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.32.0/24 | 01:05:15 | - | 22 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.33.0/24 | 01:05:15 | - | 23 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.34.0/24 | 01:05:15 | - | 42 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.35.0/24 | 01:05:15 | - | 43 | 2 | 81 | - | - | - |
| | | | | | | 80 | Yes | 24979 | - |
| B> | 100.1.36.0/24 | | - | 44 | 2 | 81 | - | - | - |

| | | | | | | | | | | |
|----|---------------|----------|---|---|--|----|-----|-------|---|---|
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 100.1.37.0/24 | 45 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 100.1.38.0/24 | 46 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 100.1.39.0/24 | 47 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 100.1.40.0/24 | 48 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1011::/64 | 24 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1012::/64 | 25 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1013::/64 | 26 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1014::/64 | 27 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1015::/64 | 28 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1016::/64 | 29 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1017::/64 | 30 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1018::/64 | 31 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1019::/64 | 32 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1020::/64 | 33 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1021::/64 | 34 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1022::/64 | 35 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |
| B> | 1023::/64 | 36 | | 2 | | 81 | - | - | - | - |
| | | 01:05:15 | - | | | 80 | Yes | 24979 | - | |
| | | 2.2.2.2 | - | | | - | | | | |

```

B> 1024::/64      37      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1025::/64      38      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1026::/64      39      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1027::/64      40      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1028::/64      41      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1029::/64      49      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1030::/64      50      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1031::/64      51      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1032::/64      52      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1033::/64      53      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1034::/64      54      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1035::/64      55      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1036::/64      56      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1037::/64      57      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1038::/64      58      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1039::/64      59      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -
B> 1140::/64      60      2      81      -      -      -      -
    01:05:15 -
        2.2.2.2      -      -      80      Yes  24979      -

```

PE4#

PE4#show ipv6 bgp summary

```
BGP router identifier 44.44.44.44, local AS number 123456
BGP table version is 9
2 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|-----|--------|--------|--------|-----|------|----------|--------------|------|
| 1641::2 | 4 | 100 | 12 | 18 | 9 | 0 | 0 | 00:03:44 | 10 | |
| 1642::2 | 4 | 100 | 11 | 16 | 9 | 0 | 0 | 00:03:45 | 10 | |
| 1643::2 | 4 | 100 | 11 | 16 | 9 | 0 | 0 | 00:03:44 | 10 | |
| 1644::2 | 4 | 100 | 11 | 16 | 9 | 0 | 0 | 00:03:44 | 10 | |
| 1645::2 | 4 | 100 | 11 | 16 | 9 | 0 | 0 | 00:03:43 | 10 | |
| 1646::2 | 4 | 100 | 11 | 16 | 9 | 0 | 0 | 00:03:42 | 10 | |
| 1647::2 | 4 | 100 | 11 | 16 | 9 | 0 | 0 | 00:03:42 | 10 | |
| 1648::2 | 4 | 100 | 11 | 16 | 9 | 0 | 0 | 00:03:44 | 10 | |
| 1649::2 | 4 | 100 | 11 | 15 | 9 | 0 | 0 | 00:03:39 | 10 | |
| 1650::2 | 4 | 100 | 11 | 15 | 9 | 0 | 0 | 00:03:40 | 10 | |

Total number of neighbors 10

```
Total number of Established sessions 10
BGP router identifier 104.1.100.1, local AS number 123456
BGP VRF l3vpn_vrf11 Route Distinguisher: 44.44.44.44:11
BGP table version is 1
2 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|-----|--------|--------|--------|-----|------|----------|--------------|------|
| 1411::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:44 | 0 | |
| 1412::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:44 | 0 | |
| 1413::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:46 | 0 | |
| 1414::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:46 | 0 | |
| 1415::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:42 | 0 | |
| 1416::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:42 | 0 | |
| 1417::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:42 | 0 | |
| 1418::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:46 | 0 | |
| 1419::2 | 4 | 100 | 10 | 14 | 1 | 0 | 0 | 00:03:44 | 0 | |
| 1420::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:42 | 0 | |
| 1421::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:44 | 0 | |
| 1422::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:44 | 0 | |
| 1423::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:46 | 0 | |
| 1424::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:42 | 0 | |
| 1425::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:44 | 0 | |
| 1426::2 | 4 | 100 | 11 | 15 | 1 | 0 | 0 | 00:03:40 | 0 | |
| 1427::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:43 | 0 | |
| 1428::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:45 | 0 | |
| 1429::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:43 | 0 | |
| 1430::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:44 | 0 | |
| 1431::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:40 | 0 | |
| 1432::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:45 | 0 | |
| 1433::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:42 | 0 | |
| 1434::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:40 | 0 | |
| 1435::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:44 | 0 | |
| 1436::2 | 4 | 100 | 10 | 14 | 1 | 0 | 0 | 00:03:45 | 0 | |
| 1437::2 | 4 | 100 | 10 | 14 | 1 | 0 | 0 | 00:03:40 | 0 | |
| 1438::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:44 | 0 | |
| 1439::2 | 4 | 100 | 10 | 14 | 1 | 0 | 0 | 00:03:45 | 0 | |
| 1440::2 | 4 | 100 | 10 | 13 | 1 | 0 | 0 | 00:03:42 | 0 | |

Total number of neighbors 30

```
Total number of Established sessions 30
PE4#
```

```
PE4#sh ip bgp vpnv4 all summary
BGP router identifier 44.44.44.44, local AS number 123456
BGP table version is 6
2 BGP AS-PATH entries
0 BGP community entries
```


| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|--------|--------|--------|--------|-----|------|----------|--------------|------|
| 2.2.2.2 | 4 | 123456 | 299 | 312 | 6 | 0 | 0 | 01:55:10 | | 60 |

Total number of neighbors 1

Total number of Established sessions 1

BGP VRF l3vpn_vrf11 Route Distinguisher: 44.44.44.44:11

BGP table version is 1

2 BGP AS-PATH entries

0 BGP community entries

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|------------|---|-----|--------|--------|--------|-----|------|----------|--------------|------|
| 104.1.11.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.12.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:57 | 10000 | |
| 104.1.13.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:59 | 10000 | |
| 104.1.14.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:59 | 10000 | |
| 104.1.15.2 | 4 | 100 | 21 | 13 | 1 | 0 | 0 | 00:04:01 | 10000 | |
| 104.1.16.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:04:01 | 10000 | |
| 104.1.17.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:57 | 10000 | |
| 104.1.18.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:57 | 10000 | |
| 104.1.19.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.20.2 | 4 | 100 | 21 | 22 | 1 | 0 | 0 | 00:03:55 | 10000 | |
| 104.1.21.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:59 | 10000 | |
| 104.1.22.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.23.2 | 4 | 100 | 21 | 22 | 1 | 0 | 0 | 00:03:57 | 10000 | |
| 104.1.24.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.25.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.26.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:55 | 10000 | |
| 104.1.27.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:57 | 10000 | |
| 104.1.28.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:56 | 10000 | |
| 104.1.29.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:57 | 10000 | |
| 104.1.30.2 | 4 | 100 | 21 | 22 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.31.2 | 4 | 100 | 21 | 22 | 1 | 0 | 0 | 00:03:57 | 10000 | |
| 104.1.32.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.33.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:59 | 10000 | |
| 104.1.34.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:57 | 10000 | |
| 104.1.35.2 | 4 | 100 | 21 | 22 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.36.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.37.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:59 | 10000 | |
| 104.1.38.2 | 4 | 100 | 21 | 22 | 1 | 0 | 0 | 00:03:57 | 10000 | |
| 104.1.39.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:58 | 10000 | |
| 104.1.40.2 | 4 | 100 | 21 | 23 | 1 | 0 | 0 | 00:03:55 | 10000 | |

Total number of neighbors 30

Total number of Established sessions 30

PE4#

New CLI

mpls l2vpn-stitching

Use this command to enable the L2VPN-IP/L3VPN stitching.

Use `no` parameter of this command to disable L2VPN-IP/L3VPN stitching.

Command Syntax

```
mpls l2vpn-stitching
```

```
no mpls l2vpn-stitching
```

Parameters

None

Default

Disabled

Command Mode

Config mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

The following example shows how to enable L2VPN-IP/L3VPN stitching using a BVI interface associated with a VPLS instance.

```
#configure terminal
(config)#mpls l2vpn-stitching
(config)#int bvi1
(config-if)#ip address 100.1.1.3/24
(config-if)#encapsulation dot1q 1001
(config-if)#access-if-vpls
(config-acc-if-vpls)#mpls-vpls vl
```

Implementation Examples

Scenario 1: A customer wants traffic originating from an L3VPN site to reach another site, but part of the service path in the provider network is delivered as L2VPN.

Use Case 1: By using L3VPN–L2VPN stitching on the same PE board, the provider can interconnect the L3VPN VRF with the L2VPN bridge-domain locally.

This avoids the need for: A second node or a physical loopback cable.

Scenario 2: A customer may have services spread across L2 (VPLS) and L3VPN domains, but they require a single, consistent default gateway location. Instead of deploying a customer router at every site, the service provider's PE must function as the central gateway for all connected locations—whether they arrive via VPLS or L3VPN.

Use Case 2: By stitching the VPLS bridge domain to the L3VPN VRF through BVI interface, the PE provides gateway functionality for both service types. L2 sites receive L3 routing without additional CE routers, and L3VPN sites reach the same gateway through normal MPLS routing. This design simplifies customer edge deployments, centralizes gateway control, and ensures consistent IP addressing and policy enforcement across mixed L2 and L3 connectivity.

Scenario 3: During network modernization, customers with existing VPLS instances may transition to L3VPN.

Some sites remain on VPLS while others move to L3VPN. Without a stitching method, the two service types cannot communicate, causing service disruption during migration.

Use Case 3: By introducing BVI interface between the VPLS bridge domain and the L3VPN VRF, sites on both services can communicate during the transition period. This ensures gradual migration, avoids downtime, and minimizes configuration changes across the network.

Troubleshooting

1. BVI IP is not responding to ping.

- Check that L2VPN-STITCHING feature is enabled using "show hardware-profile filters" command. L2VPN-STITCHING group should be present.
- Check BVI interface status. It should be up in "show mpls vpls" and "show ip int br" command.
- Check if incoming pwe-label is properly mapped to bvi using "show hsl mpls l2vpn bvi-label-tree" command.
- Check if incoming traffic has same encapsulation as BVI interface.

2. Traffic is not routed out from BVI interface.

- Check that L2VPN-STITCHING feature is enabled using "show hardware-profile filters" command. L2VPN-STITCHING group should be present.
- Check BVI interface status. It should be up in "show mpls vpls" and "show ip int br" command.
- FOR Q1 devices, check if proper oper count is maintained (1 per ac and 1 per pwe for the vpls).
- FOR Q2 devices, check if bvi and destination phy (ac or pwe port) is present in "show hsl hw unit 0 vpn-ports <BVI_VID>"
- Check if icmpv6|arp|dhcp flags are enabled on ac and pwe lif using "show hsl hw unit 0 vpn-ports <BVI_VID>" and "show hsl hw unit 0 pwe-lif" command.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms/Acronym | Description |
|-------------------|---------------------------------|
| BVI | Bridged Virtual Interface |
| FEC | Forwarding Equivalence Class |
| IRB | Integrated Routing and Bridging |
| L2VPN | Layer 2 Virtual Private Network |
| L3VPN | Layer 3 Virtual Private Network |
| LDP | Label Distribution Protocol |
| VPLS | Virtual Private LAN Service |
| VPWS | Virtual Private Wire Service |

MULTI-PROTOCOL LABEL SWITCHING

COMMAND REFERENCE

| | |
|--|------|
| MPLS Commands | 1381 |
| admin-groups | 1384 |
| advanced parser enable | 1385 |
| bandwidth | 1386 |
| clear mpls counters ldp | 1387 |
| clear mpls counters rsvp | 1388 |
| clear mpls counters static | 1389 |
| clear mpls l2-circuit statistics | 1390 |
| control-word | 1391 |
| flow-label | 1392 |
| group-id | 1394 |
| group-name | 1395 |
| label-switching | 1396 |
| manual-pseudowire | 1397 |
| match vlan | 1398 |
| mpls ac-group | 1400 |
| mpls bandwidth-class | 1401 |
| mpls-bos parsing nibble-map | 1402 |
| mpls ftn-ecmp ldp | 1403 |
| mpls ftn-entry tunnel-id | 1404 |
| mpls ftn-entry | 1406 |
| mpls ilm-ecmp ldp | 1407 |
| mpls ilm-entry pop | 1408 |
| mpls ilm-entry swap | 1409 |
| mpls ingress-ttl | 1411 |
| mpls l2-circuit | 1412 |
| mpls l2-circuit-fib-entry | 1413 |
| mpls-l2vpn-nw-statistics | 1414 |
| mpls label mode | 1415 |
| mpls local-packet-handling | 1417 |
| mpls lsp-encap-dscp-preserve | 1418 |
| mpls lsp-model uniform | 1419 |
| mpls lsp-stitching | 1420 |
| mpls map-route | 1421 |
| mpls min-label-value | 1422 |

| | |
|--|------|
| mpls propagate-ttl | 1424 |
| mpls traffic-eng srlg | 1425 |
| pce entity | 1426 |
| pce lsp-delegate | 1427 |
| pce state-report | 1428 |
| ping mpls | 1429 |
| rewrite ingress | 1433 |
| secondary-priority srlg-disjoint | 1434 |
| secondary srlg-disjoint | 1435 |
| service-tpid | 1436 |
| show admin-groups | 1437 |
| show hsl mpls l3vpn-ftn | 1438 |
| show mpls bandwidth-class | 1439 |
| show mpls counters ldp | 1440 |
| show mpls | 1442 |
| show mpls counters static | 1444 |
| show mpls counters rsvp | 1446 |
| show mpls cross-connect-table | 1448 |
| show mpls cross-connect-table backup | 1450 |
| show mpls forwarding-table detail | 1451 |
| show mpls forwarding-table | 1453 |
| show mpls ftn-table | 1456 |
| show mpls ilm-table | 1458 |
| show mpls in-segment-table | 1461 |
| show mpls l2-circuit | 1463 |
| show mpls l2-circuit statistics | 1465 |
| show mpls mapped-routes | 1467 |
| show mpls out-segment-table | 1468 |
| show mpls owner-rib-table | 1470 |
| show mpls qos-resource | 1471 |
| show mpls vc-table | 1473 |
| show mpls vrf | 1474 |
| show mpls vrf-forwarding-table vrf | 1475 |
| show running-config interface mpls | 1478 |
| show running-config mpls | 1479 |
| show running-config vc | 1480 |
| show running-config vpls | 1481 |
| show vccv statistics | 1482 |
| srlg-disjoint | 1483 |
| trace mpls | 1485 |

| | |
|--|------|
| suppress-oper-log mpls l2vpn | 1487 |
| trace mpls | 1488 |
| tunnel-name | 1490 |
| tunnel-id | 1491 |
| tunnel-select-policy | 1492 |
| vccv cc-type | 1493 |
| vccv cv-type | 1494 |
| Virtual Private LAN Service Commands | 1495 |
| ac-admin-status | 1496 |
| ac-description | 1497 |
| clear mpls vpls | 1498 |
| control-word | 1499 |
| exit-signaling | 1500 |
| exit-if-vpls | 1501 |
| learning disable (VPLS Mode) | 1502 |
| learning disable (Interface VPLS Mode) | 1503 |
| learning enable | 1504 |
| mpls vpls | 1505 |
| mpls-vpls service-template | 1506 |
| no learning | 1507 |
| show mpls vpls | 1508 |
| show mpls vpls mac-address | 1516 |
| show mpls vpls split-horizon | 1518 |
| show mpls vpls statistics | 1520 |
| signaling ldp | 1522 |
| signaling bgp | 1523 |
| static-mac | 1524 |
| tunnel-select-policy | 1525 |
| ve-id | 1526 |
| vpls-ac-group | 1527 |
| vpls-description | 1528 |
| vpls fib-entry | 1529 |
| vpls-mtu | 1531 |
| vpls-peer | 1532 |
| vpls-peer manual | 1534 |
| vpls-type | 1535 |
| vpls-vc | 1536 |
| vpls-route-map | 1537 |
| EVPN MPLS Commands | 1538 |
| access-if-evpn | 1540 |

| | |
|--|------|
| arp-cache disable | 1541 |
| arp-nd cos | 1542 |
| arp-nd flood-suppress | 1543 |
| arp-nd refresh timer | 1544 |
| clear evpn mpls counters | 1545 |
| clear mac address-table | 1546 |
| dynamic-learning disable | 1547 |
| evi-name | 1548 |
| evpn esi hold-time | 1549 |
| evpn etree | 1550 |
| evpn irb | 1551 |
| evpn irb-forwarding anycast-gateway-mac | 1552 |
| evpn irb-if-forwarding anycast-gateway-mac | 1553 |
| evpn mpls enable | 1554 |
| evpn mpls id | 1555 |
| evpn mpls irb | 1557 |
| l3vni | 1558 |
| mac | 1559 |
| evpn mpls mac-ageing-time | 1560 |
| evpn mpls multihoming enable | 1561 |
| evpn mpls vtep-ip-global | 1562 |
| evpn multi-homed | 1563 |
| hardware-profile filter evpn-mpls-cw | 1565 |
| hardware-profile filter evpn-mpls-mh | 1566 |
| hardware-profile statistics evpn-mpls | 1567 |
| host-reachability-protocol | 1568 |
| evpn-route-map | 1569 |
| ip address | 1570 |
| ipv6 address | 1571 |
| link-loss forwarding | 1572 |
| mac-holdtime | 1573 |
| mac-vrf | 1574 |
| map vpn-id | 1575 |
| nd-cache disable | 1576 |
| rewrite egr-pcp-disable | 1577 |
| service-carving | 1578 |
| service-carving ac-driven | 1580 |
| service-carving weight | 1581 |
| show bgp l2vpn evpn | 1582 |
| show evpn load-balance | 1584 |

| | |
|--|------|
| show evpn mpls | 1586 |
| show evpn mpls arp-cache | 1587 |
| show evpn mpls counters | 1588 |
| show evpn mpls label alias | 1589 |
| show evpn mpls label esi | 1590 |
| show evpn mpls mac-table | 1591 |
| show evpn mpls nd-cache | 1593 |
| show evpn mpls route-count | 1594 |
| show evpn mpls static host state | 1596 |
| show evpn mpls status | 1597 |
| show evpn mpls tunnel | 1598 |
| show evpn mpls tunnel label | 1599 |
| show evpn mpls xconnect | 1600 |
| show evpn mpls xconnect tunnel | 1601 |
| show evpn mpls xconnect tunnel label | 1602 |
| show running-config evpn mpls | 1603 |
| tunnel-select-policy | 1605 |
| MPLS Service Mapping Commands | 1606 |
| color | 1607 |
| mode | 1608 |
| mpls 6pe tunnel-select-policy | 1609 |
| show running-config tunnel-policy | 1610 |
| show tunnel-policy | 1611 |
| tunnel-policy | 1614 |
| tunnel-select-policy (For L3VPN service) | 1615 |
| MPLS LSP BFD Commands | 1616 |
| lsp ping-interval | 1617 |
| min-tx | 1618 |
| mpls bfd | 1619 |
| shutdown | 1620 |
| L2VPN over L2 Subinterface Commands | 1621 |
| access-if-vpws | 1622 |
| mpls-l2-circuit | 1623 |
| vc-mode | 1624 |
| access-if-vpls | 1625 |
| mpls-vpls | 1626 |
| learning disable | 1627 |
| split-horizon (subinterface) | 1628 |
| static-mac | 1629 |

MPLS Commands

This chapter is a reference for the MPLS commands:

| | |
|--|------|
| admin-groups | 1384 |
| advanced parser enable | 1385 |
| bandwidth | 1386 |
| clear mpls counters ldp | 1387 |
| clear mpls counters rsvp | 1388 |
| clear mpls counters static | 1389 |
| clear mpls l2-circuit statistics | 1390 |
| control-word | 1391 |
| flow-label | 1392 |
| group-id | 1394 |
| group-name | 1395 |
| label-switching | 1396 |
| manual-pseudowire | 1397 |
| match vlan | 1398 |
| mpls ac-group | 1400 |
| mpls bandwidth-class | 1401 |
| mpls-bos parsing nibble-map | 1402 |
| mpls ftn-ecmp ldp | 1403 |
| mpls ftn-entry tunnel-id | 1404 |
| mpls ftn-entry | 1406 |
| mpls ilm-ecmp ldp | 1407 |
| mpls ilm-entry pop | 1408 |
| mpls ilm-entry swap | 1409 |
| mpls ingress-ttl | 1411 |
| mpls l2-circuit | 1412 |
| mpls l2-circuit-fib-entry | 1413 |
| mpls-l2vpn-nw-statistics | 1414 |
| mpls label mode | 1415 |
| mpls local-packet-handling | 1417 |
| mpls lsp-encap-dscp-preserve | 1418 |
| mpls lsp-model uniform | 1419 |
| mpls lsp-stitching | 1420 |
| mpls map-route | 1421 |
| mpls min-label-value | 1422 |
| mpls propagate-ttl | 1424 |
| mpls traffic-eng srlg | 1425 |

| | |
|--|------|
| pce entity | 1426 |
| pce lsp-delegate | 1427 |
| pce state-report | 1428 |
| ping mpls | 1429 |
| rewrite ingress | 1433 |
| secondary-priority srlg-disjoint | 1434 |
| secondary srlg-disjoint | 1435 |
| service-tpid | 1436 |
| show admin-groups | 1437 |
| show hsl mpls l3vpn-fts | 1438 |
| show mpls bandwidth-class | 1439 |
| show mpls counters ldp | 1440 |
| show mpls | 1442 |
| show mpls counters static | 1444 |
| show mpls counters rsvp | 1446 |
| show mpls cross-connect-table | 1448 |
| show mpls cross-connect-table backup | 1450 |
| show mpls forwarding-table detail | 1451 |
| show mpls forwarding-table | 1453 |
| show mpls ftn-table | 1456 |
| show mpls ilm-table | 1458 |
| show mpls in-segment-table | 1461 |
| show mpls l2-circuit | 1463 |
| show mpls l2-circuit statistics | 1465 |
| show mpls mapped-routes | 1467 |
| show mpls out-segment-table | 1468 |
| show mpls owner-rib-table | 1470 |
| show mpls qos-resource | 1471 |
| show mpls vc-table | 1473 |
| show mpls vrf | 1474 |
| show mpls vrf-forwarding-table vrf | 1475 |
| show running-config interface mpls | 1478 |
| show running-config mpls | 1479 |
| show running-config vc | 1480 |
| show running-config vpls | 1481 |
| show vccv statistics | 1482 |
| srlg-disjoint | 1483 |
| trace mpls | 1485 |
| suppress-oper-log mpls l2vpn | 1487 |
| trace mpls | 1488 |

| | |
|----------------------------|------|
| tunnel-name | 1490 |
| tunnel-id | 1491 |
| tunnel-select-policy | 1492 |
| vccv cc-type | 1493 |
| vccv cv-type | 1494 |

admin-groups

Use this command to create a name-to-value binding for an administrative group.



Note: Only 32 administrative groups can be configured at one time.

Use the `no` parameter with this command to remove a named administrative group.

Command Syntax

```
admin-group NAME <0-31>  
no admin-group NAME <0-31>
```

Parameters

NAME

Name of administrative group

<0-31>

The value of the administrative group

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#admin-group mygroup 3
```

advanced parser enable

Use this command to enable the DPI on Qumran1 (Q1) series platforms.

Command Syntax

```
advanced parser enable
```

Parameters

None

Default

None

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.3 and applicable to Qumran1 (Q1) series platforms.

Example

This example shows how to enable the DPI on Qumran1 (Q1) series platforms:

```
OcNOS#configure terminal
OcNOS#(config)#advanced parser enable
OcNOS#(config)#commit
```

bandwidth

Use this command to specify the maximum bandwidth to be used for a band-class. The bandwidth value is in bits.



Note: Run this command in the Bandwidth-class mode (refer to [mpls bandwidth-class \(page 1401\)](#)).

Command Syntax

```
bandwidth BANDWIDTH setup-priority <0-7> hold-priority <0-7>
```

Parameter

BANDWIDTH

<1-999>k for 1 to 999 kilo bits/s

<1-999>m for 1 to 999 mega bits/s

<1-100>g for 1 to 100 giga bits/s

setup-priority

Indicate the `setup-priority` parameter

<0-7>

The actual setup priority value

hold-priority

Indicate the `hold-priority` parameter

<0-7>

The actual hold priority value

Default

Zero

Command Mode

Bandwidth-class mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#mpls bandwidth-class new-BC
(config-mpls-bw)#bandwidth 100m setup-priority 1 hold-priority 1
```

clear mpls counters ldp

Use this command to clear traffic statistics for FTNs and ILMs configured by LDP.

Command Syntax

```
clear mpls counters ldp ((ftn (|A.B.C.D/M)) | (ilm (|A.B.C.D/M)) |)
```

Parameter

ftn

FEC-to-NHLFE map counters

A.B.C.D/M

FEC prefix

ilm

Incoming label map counters

A.B.C.D/M

FEC prefix

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear mpls counters ldp
```

clear mpls counters rsvp

Use this command to clear traffic statistics for LSPs configured by RSVP.

Command Syntax

```
clear mpls counters rsvp ((tunnel-name NAME) | (tunnel-id TUNNEL_ID) | (node-role (ingress | transit | egress)) |)
```

Parameter

NAME

RSVP tunnel name

TUNNEL_ID

RSVP tunnel identifier

ingress

LSP role is ingress

transit

LSP role is transit

egress

LSP role is egress

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear mpls counters rsvp
```


clear mpls counters static

Use this command to clear traffic statistics for statically configured FTNs and ILMs.

Command Syntax

```
clear mpls counters static ((ftn (|A.B.C.D/M)) | (ilm (|A.B.C.D/M)) |)
```

Parameter

ftn

FEC-to-NHLFE map counters

A.B.C.D/M

FEC prefix

ilm

Incoming label map counters

A.B.C.D/M

FEC prefix

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#clear mpls counters static
```

clear mpls l2-circuit statistics

Use this command to clear MPLS traffic statistics for L2 circuit.

Command Syntax

```
clear mpls l2-circuit NAME statistics (access-port|network-port|)
```

Parameters

name

Name of L2 circuit

access-port

Displays the access port statistics

network-port

Displays the network port statistics

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear mpls l2-circuit vcl statistics
```

control-word

Use this command to enable control word for the MPLS layer-2 virtual circuit.

Use the no parameter with this command to disable control word from the MPLS layer-2 virtual circuit.



Notes:

- Control-word (CW) negotiation can be `Preferred` or `Not Preferred`. OcNOS follows the `Not Preferred` mode. OcNOS does not rely on whether the peer has control-word enabled.
 - If both sides enable control-word, OcNOS installs the negotiated CW as 1.
 - If only one side enables control-word, OcNOS installs the negotiated CW as 0.
 - In either case, OcNOS does not bring the Virtual Circuit (VC) down.
- Some interoperability devices use the `Preferred` mode in control-word negotiation. When these devices are configured with control-word, they expect the other end device (OcNOS) to also have control-word configured. If the OcNOS does not match its configuration, the interoperability devices set the PW status to `Not Forwarding` or raise a fault. As a result, the VC on OcNOS appears as `Down` because the remote pseudowire (PW) status indicates a fault.

Command Syntax

```
control-word
no control-word
```

Parameters

None

Default

Disabled

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4
(config-pseudowire)#control-word
```

flow-label

Use this command to enable flow-label transmit, receive or both for the MPLS layer-2 virtual circuit and mpls vpls on dynamic and static VPWS and VPLS.

Use `no` command to disable flow-label transmit, receive or both for the MPLS layer-2 virtual circuit and mpls vpls on dynamic and static VPWS and VPLS.

For signaling flow-label capability in an RFC-compliant (RFC6391/8077) way using "The PWid FEC Element" = 0x80 when all Interface Parameter Sub-TLV is part of PWid FEC Element to support interop with other vendors, then interface-param-tlv can be used.

Command Syntax

```
flow-label (both|receive|transmit) (static|interface-param-tlv|)
no flow-label (static|interface-param-tlv|)
```

Parameters

both

flow label direction both (transmit and receive)

receive

flow label direction receive

transmit

flow label direction transmit

static

flow-label static

interface-param-tlv

ldp specific interface parameter tlv

Command Mode

config-pseudowire mode and config-vpls mode

Applicability

This command was introduced in OcNOS version 6.0.0 and added parameter `interface-param-tlv` in OcNOS version 6.3.0 to support interop with other vendors.

Example

For the config-pseudowire mode:

```
# configure terminal
(config)#mpls l2-circuit vc1111 1111 7.7.7.7
(config-pseudowire)#flow-label transmit

# configure terminal
(config)#mpls l2-circuit vc1111 1111 7.7.7.7
(config-pseudowire)#no flow-label
```

For the config-vpls mode:

```
# configure terminal
(config)#mpls vpls test 100
(config-vpls)# flow-label both interface-param-tlv
# configure terminal
(config)#mpls vpls test 100
(config-vpls)# no flow-label
```

For the static configuration:

```
# configure terminal
(config)#mpls l2-circuit vc1111 1111 7.7.7.7
(config-pseudowire)#flow-label transmit static
```

For the config-vpls mode:

```
# configure terminal
(config)#mpls vpls test 100
(config-vpls)# flow-label both static

# configure terminal
(config)#mpls vpls test 100
(config-vpls)# no flow-label static
```

For the un-configuration of the static flow label:

For the config-pseudowire mode:

```
# configure terminal
(config)#mpls l2-circuit vc1111 1111 7.7.7.7
(config-pseudowire)#no flow-label static
```

group-id

Use this command to configure a specific group identifier to existing group with a group name in the MPLS layer-2 virtual circuit.

Use the no parameter with this command to remove group identifier from the MPLS layer-2 virtual circuit

Command Syntax

```
group-id <1-4294967295>  
no group-id
```

Parameters

<1-4294967295>

Value for group identifier

Default

By default, group-id is disabled. If group-name is configured, default group-id is the first available identifier.

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4  
(config-pseudowire)#group-name group-1  
(config-pseudowire)#group-id 11
```

group-name

Use this command to map the MPLS layer-2 virtual circuit with a specific group.

Use the no parameter with this command to remove group from the MPLS layer-2 virtual circuit

Command Syntax

```
group-name NAME  
no group-name
```

Parameters

NAME

String identifying group NAME

Default

Disabled

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 6.0.0.

Example

```
#configure terminal  
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4  
(config-pseudowire)#group-name group-1
```

label-switching

Use this command to either enable label-switching on an interface or to modify the label-space to which this interface is bound.

Use the `no` parameter and the interface is bound to the platform-wide (zero) label-space.



Notes:

- When label-switching enabled on VLAN interface, MTU value must be manually increased by at least 20 bytes on Parent interfaces of VLAN. Example, default MTU must be set as 1520 instead of 1500 on label-switching parent interface label switched VLAN interface. (Parent Interface MTU \geq label switched VLAN interface MTU + 20).
- Due to hardware solution restriction, MPLS termination on double-tagged sub-interface is not supported for Qumran1 (Q1) series platforms.

Command Syntax

```
label-switching
label-switching <0-60000>
no label-switching
```

Parameter

<0-60000>

Label space value in this range

Default

Disabled

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows the enabling of label switching on the `eth0` interface.

```
#configure terminal
(config)#interface eth0
(config-if)#label-switching 654
```


manual-pseudowire

Use this command to configure the MPLS layer-2 virtual circuit as manual. This disables the default signalling mode of the MPLS layer-2 virtual circuit.

Use the no parameter with this command to disable manual mode of the MPLS layer-2 virtual circuit & enable signalling mode of the MPLS layer-2 virtual circuit.

Command Syntax

```
manual-pseudowire  
no manual-pseudowire
```

Parameters

NA

Default

Disabled

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4  
(config-pseudowire)#manual-pseudowire
```

match vlan

Use this command to configure a match VLAN action for a service template.

Use the `no` parameter to remove a match VLAN action for a service template.

Command Syntax

```
match (all | double-tag outer-vlan <2-4094> inner-vlan VLAN_RANGE | outer-vlan VLAN_RANGE | untagged)
no match (double-tag outer-vlan <2-4094> inner-vlan VLAN_RANGE | outer-vlan VLAN_RANGE | untagged)
```

Parameter

all

Accept all matches

double-tag

Double tag match

outer-vlan

Double tag outer VLAN

<2-4094>

Outer VLAN identifier

inner-vlan

Double tag inner VLAN

VLAN_RANGE

VLAN identifier <2-4094> range: 2-5,10 or 2-5,7-19

outer-vlan

Single tag outer-VLAN

VLAN_RANGE

VLAN identifier <2-4094> range: 2-5,10 or 2-5,7-19

untagged

Match untagged. This parameter depends on the `switchport dot1q ethertype` configuration. Packets received with a TPID other than 0x8100 (default value) and the TPID value configured by `switchport dot1q ethertype` are treated as untagged. For example, if you give the command: `switchport dot1q ethertype 0x8888` then packets received with TPID 0x8100 or 0x88a8 are treated as tagged. Packets received with other TPIDs are treated as untagged.

Command Mode

MPLS SVC mode

Applicability

This command was introduced before OcNOS version 1.3 and updated in OcNOS version 3.0.

The inner vlan range option added in OcNOS version 4.1.

Example

```
#configure terminal
```

```
(config)#service-template C2
(config-svc)#match double-tag outer-vlan 10 inner-vlan 20
(config-svc)#exit
(config)#service-template C2
(config-svc)#no match double-tag outer-vlan 10 inner-vlan 20
(config-svc)#exit
#configure terminal
(config)#service-template C3
(config-svc)#match double-tag outer-vlan 10 inner-vlan 200-300
(config-svc)#exit
(config)#service-template C4
(config-svc)#no match double-tag outer-vlan 10 inner-vlan 200-300
(config-svc)#exit
#configure terminal
(config)#service-template t1
(config-svc)#match untagged
(config-svc)#rewrite ingress push 100
```

mpls ac-group

Use this command to create a new access circuit group for MPLS.

Use the `no` parameter with this command to remove an access circuit group.

Command Syntax

```
mpls ac-group NAME <1-4294967295>  
no mpls ac-group NAME
```

Parameter

NAME

The name of the access circuit group

<1-4294967295>

The identifier for the group; used in LDP

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#mpls ac-group new-ac 123  
  
(config)#no mpls ac-group new-ac
```

mpls bandwidth-class

Use this command to create a new bandwidth class name. Using this command changes the command mode to Bandwidth-class mode.

Use the `no` parameter with this command to remove a bandwidth class name.

Command Syntax

```
mpls bandwidth-class NAME  
no mpls bandwidth-class NAME
```

Parameter

NAME

Name of the bandwidth class

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#mpls bandwidth-class new-BC  
(config-mpls-bw)#  
  
(config)#no mpls bandwidth-class new-BC
```

mpls-bos parsing nibble-map

Use this command to configure the next protocol speculation after MPLS-BOS based on the first nibble value of zero.

The control word, legacy protocols (like Frame Relay over MPLS, ATM over MPLS, HDLC/PPP over MPLS), or an Ethernet packet can have a nibble value of 0 for its first nibble value after MPLS BOS. By default, the first nibble value of zero is parsed as an Ethernet packet in hardware. Configure the first nibble value of zero to be parsed as a control word if there is an expectation to receive a control word packet. The control word defines the underlying packet as an Ethernet packet. The legacy protocol packets cannot be parsed by hardware, so a disable option is provided.

Command Syntax

```
mpls-bos parsing nibble-map 0 (disable|cw)
```

Parameters

disable

Maps the nibble to no speculation.

cw

Maps the nibble to control-word.

Default

None

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.6.1 and applicable to Qumran1 (Q1) series platforms.

Example

This example shows how to configure the next protocol speculation after MPLS-BOS based on the first nibble values on Qumran1 (Q1) series platforms:

```
OcNOS#configure terminal
OcNOS(config)#mpls-bos parsing nibble-map 0 disable
OcNOS(config)#commit
```

mpls ftn-ecmp ldp

Use this command to enable Equal-Cost Multi-Path ECMP for Label Distribution Protocol (LDP) Forwarding Table Entry (FTN). FTN contains the details of forwarding the labeled packets.

Use `no` command to disable ECMP for LDP FTN.

Command Syntax

```
mpls ftn-ecmp ldpno mpls ftn-ecmp ldp
```

Parameter

None

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.4.1.

Examples

The below example shows how to enable and disable ECMP for LDP FTN:

```
#configure terminal
(config)#mpls ftn-ecmp ldp
(config)# commit
(config)# no mpls ftn-ecmp ldp
(config)# commit
```

mpls ftn-entry tunnel-id

This command will be used to create a static tunnel.

In hardware, it creates a logical interface to which services can be mapped.



Note: Only global label space is supported and per interface label space is not supported. MPLS core with IPv6 is not supported.

Command Syntax

```
mpls ftn-entry tunnel-id <1-5000> (A.B.C.D/M|A.B.C.D A.B.C.D) <16-1048575> A.B.C.D IFNAME
((secondary|primary)|)
no mpls ftn-entry tunnel-id <1-5000> (A.B.C.D/M|A.B.C.D A.B.C.D) <16-1048575> A.B.C.D IFNAME
((secondary|primary)|)
mpls ftn-entry tunnel-id <1-5000> X:X::X:X/M <16-1048575> X:X::X:X IFNAME ((secondary|primary)|)
no mpls ftn-entry tunnel-id <1-5000> X:X::X:X/M <16-1048575> X:X::X:X IFNAME ((secondary|primary)|)
```

Parameters

<1-5000>

The tunnel ID value

A.B.C.D/M

Forwarding equivalence class with mask

A.B.C.D

Mask for forwarding equivalency class

<16-1048575>

Outgoing label

A.B.C.D

Nexthop IPv4 address

IFNAME

Outgoing interface name

primary

The primary LSP; default is primary

secondary

The secondary LSP Command Mode

X:X::X:X/M

IPv6 Forwarding Equivalence Class with Mask

X:X::X:X

Nexthop IPv6 address

Default

Disabled

Command mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#mpls ftn-entry tunnel-id 2 10.10.0.0/24 16 1.2.3.4 eth1 secondary
(config)#no mpls ftn-entry tunnel-id 2 10.10.0.0/24 16 1.2.3.4 eth1 secondary
```

mpls ftn-entry

Use this command to create a static LSP. In the hardware, this command creates an IP route with outgoing MPLS parameters.



Note: Only global label space is supported and per interface label space is not supported.

Command Syntax

```
mpls ftn-entry (A.B.C.D/M|A.B.C.D A.B.C.D) <16-1048575> A.B.C.D IFNAME
no mpls ftn-entry (A.B.C.D/M|A.B.C.D A.B.C.D) <16-1048575> A.B.C.D IFNAME
```

Parameters

A.D.C.D/M

Forwarding Equivalence Class with Mask

A.B.C.D

Mask for forwarding equivalency class

<16-1048575>

Outgoing label <16-1048575>

A.B.C.D

Nexthop IPv4 address

IFNAME

Outgoing interface name

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)# mpls ftn-entry 2.2.2.2/32 111 20.0.0.2 eth1
(config)# no mpls ftn-entry 2.2.2.2/32 111 20.0.0.2 eth1
```

mpls ilm-ecmp ldp

Use this to enable ECMP for LDP Incoming Label Map (ILM).

Use `no` of this to disable ECMP for LDP ILM.



Notes:

- Entropy is also required to be configured for load-balancing to work.
- LDP has to be configured with `no-php` for entropy to work (Q1 platforms).

Command Syntax

```
#configure terminal
(config)# mpls ftn-entry 2.2.2.2/32 111 20.0.0.2 eth1
(config)# no mpls ftn-entry 2.2.2.2/32 111 20.0.0.2 eth1
```

Parameters

None

Default

LDP ECMP on transit nodes is disabled. If LDP load-balancing is required on transit nodes, enable this option.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 6.1.0.

Example

```
#configure terminal
(config)#mpls ilm-ecmp ldp
(config)#no mpls ilm-ecmp ldp
```

mpls ilm-entry pop

Use this command to create an ILM entry in the ILM table to which a POP incoming interface is bound. Upon receipt of a labeled packet on an MPLS-enabled router, a lookup is done based on the incoming label in the ILM table. If a match is found, the packet may either be label-switched downstream, or popped and passed over IP. In a pop operation, an outgoing label is not needed as is either accepted or forwarded over IP. The nexthop option is also not mandatory because the FEC IP address could be a local IP address.

Use the `no` option with the command to delete an ILM entry. If there is no match, an error message displays.



Note: Only global label space is supported and per interface label space is not supported.

Command Syntax

```
mpls ilm-entry <16-1048575> pop
no mpls ilm-entry <16-1048575> pop
```

Parameters

<16-1048575>

Incoming label value

pop

Pop the incoming label

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#mpls ilm-entry 100 pop
```

mpls ilm-entry swap

Use this command to create an ILM entry in the ILM table to which a swap incoming interface is bound. Upon receipt of a labeled packet on an MPLS-enabled router, a lookup is done based on the incoming label in the ILM table. If a match is found, the packet may either be label-switched downstream, or popped and passed over IP.

Use the `no` option with the command to delete an ILM entry. If there is no match, an error message displays.



Note: Only global label space is supported and per interface label space is not supported. MPLS core with IPv6 is not supported.

Command Syntax

```
mpls ilm-entry <16-1048575> swap <16-1048575> IFNAME A.B.C.D (A.B.C.D/M|A.B.C.D A.B.C.D)
no mpls ilm-entry <16-1048575> swap <16-1048575> IFNAME A.B.C.D (A.B.C.D/M|A.B.C.D A.B.C.D)
mpls ilm-entry <16-1048575> swap <16-1048575> IFNAME X:X::X:X X:X::X:X/M
no mpls ilm-entry <16-1048575> swap <16-1048575> IFNAME X:X::X:X X:X::X:X/M
```

Parameters

<16-1048575>

Incoming label value range <16-1048575>

swap

Specify swap for the incoming label

<16-1048575>

Configure an outgoing label with a value from <16-1048575>



Note: A value of 2 indicates explicit NULL and a value of 3 indicates implicit NULL.

IFNAME

Outgoing interface name

A.B.C.D

Nexthop IPv4 address

A.B.C.D

The FEC for which this ILM entry is created

A.B.C.D/M

The FEC for which this ILM entry is created, plus mask

A.B.C.D

A mask for forwarding equivalence class mask

X:X::X:X/M

IPv6 FEC for which this ILM entry is being created, plus mask

X:X::X:X

Nexthop IPv6 address

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3 and the incoming interface name parameter `IFNAME` is removed in OcNOS version 6.4.1.

Examples

```
#configure terminal
(config)#mpls ilm-entry 16 swap 17 eth2 1.1.1.1 1.1.1.1/3
```

mpls ingress-ttl

Use this command to set a Time to Live (TTL) value for LSPs for which this LSR is the ingress.

Use the `no` parameter with this command to unset the custom TTL value being used for LSPs for which this LSR is the ingress.

Command Syntax

```
mpls ingress-ttl <0-255>  
no mpls ingress-ttl
```

Parameter

<0-255>

Set the TTL value to use

Default

64

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#mpls ingress-ttl 3
```

mpls l2-circuit

Use this command to create an instance of an MPLS layer 2 virtual circuit, without specifying a group to which the VC belongs. Refer to [group-name \(page 1395\)](#) for information on how to create an MPLS “with” a specific group. A Layer-2 MPLS Virtual Circuit instance may be bound to any interface on the router; however, only one interface may be bound to a Layer-2 circuit at a time.

Use the `no` parameter with this command to delete an instance of an MPLS Layer-2 Virtual Circuit.



Note: OcNOS supports configuring the same Virtual Circuit Identifier (VC-ID) with different VPWS names for different LDP neighbors.

Command Syntax

```
mpls l2-circuit NAME <1-4294967295> A.B.C.D
mpls l2-circuit NAME <1-4294967295> A.B.C.D mode raw
mpls l2-circuit NAME <1-4294967295> A.B.C.D mode tagged
no mpls l2-circuit NAME <1-4294967295> A.B.C.D
```

Parameters

NAME

String identifying the MPLS Layer-2 virtual circuit

<1-4294967295>

A 32-bit identifier to which the L2 circuit name should be mapped

A.B.C.D

IPv4 address for the MPLS L2 virtual circuit end-point

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4
```


mpls l2-circuit-fib-entry

Use this command to add a static Layer-2 MPLS Virtual Circuit FIB entry.

Use the no parameter with this command to delete a Layer-2 MPLS Virtual Circuit FIB entry.

Command Syntax

```
mpls l2-circuit-fib-entry <1-4294967295> <16-15999> <16-15999> (A.B.C.D|X:X::X:X) IFNAME NAME
no mpls l2-circuit-fib-entry <1-4294967295> (A.B.C.D|X:X::X:X)
```

Parameters

fib-entry <1-4294967295>

The Virtual Circuit ID ranges from 1 to 4294967295

fib-entry <16-15999>

Incoming label in the range of <16-1048575>

fib-entry <16-15999>

Outgoing label in the range of <16-1048585>

fib-entry A.B.C.D

Nexthop IPv4 address

fib-entry X:X::X:X

Nexthop IPv6 address

fib-entry IFNAME

Provider-facing interface name

fib-entry NAME

Name of the access interface or the connecting VC.

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#mpls l2-circuit-fib-entry 10 100 200 10.10.10.10 eth1 eth2

#configure terminal
(config)#no mpls l2-circuit-fib-entry 10 10.10.10.10
```

mpls-l2vpn-nw-statistics

Use this command to enable traffic statistics collection for L2VPN services (VPLS and VPWS) on Qumran2 (Q2) series platforms.

On Qumran2 (Q2) series platforms, a single LIF-based statistics profile is used across the entire solution. As a result, the L2VPN traffic statistics can be collected either at L2VPN service or at transport, but not both simultaneously. By default, traffic statistics are collected at transport level.

If users prefer to collect L2VPN statistics, they can enable this command at the cost of losing transport-level statistics.



Note: When the auto bandwidth feature is deployed, the `mpls-l2vpn-nw-statistics` command must not be configured, as it will affect the accuracy of the bandwidth calculation.

Use the `no` form of this parameter to disable traffic statistics collection at L2VPN service level.

Command Syntax

```
mpls l2vpn-nw-statistics
no mpls l2vpn-nw-statistics
```

Parameters

None

Command Mode

Configuration mode

Default

Disabled

Applicability

Introduced in OcNOS version 6.6.1 and applicable only for Qumran2 (Q2) series platforms.

Example

```
#config terminal
(config)#mpls l2vpn-nw-statistics
(config)#commit

(config)#no mpls l2vpn-nw-statistics
```

mpls label mode

Use this command to configure label allocation mode for VPNv4 and/or VPNv6 routes. Label allocation mode as per-vrf is the default mode in which single mpls-label is allocated for all VPN Routes in a VRF. Label allocation mode as per-prefix will allocate unique mpls-labels per VPN route in a VRF. If allocation model is disabled using no mpls label mode configuration, the configuration reverts back to default-mode .

Label allocation mode is the local property i.e. the VRF routes are distributed to BGP-peer as per the mode configured on local node. When per-vrf mode is configured, single label for all routes in the VRF will be distributed to peer node.

Label allocation mode can be set for all VRFs or selective VRFs by these commands:

```
mpls label mode vpnv4 all-vrfs per-vrf
```

- If the admin selects the per-vrf mode for the entire system, then all VRFs switches to per-vrf allocation mode except for the VRFs that has been explicitly configured using command mpls label mode vpnv4 vrf WORD per-prefix. Label allocation mode set using specific VRF takes precedence over all-vrf command.

```
mpls label mode vpnv6 vrf WORD per-vrf
```

- If the admin selects per-vrf mode for a particular vrf say vrf1, then only vrf1 switches to per-vrf mode and rest of the vrfs will remain in default allocation mode.



Note: The CLI requires an explicit `clear ip bgp` command to observe the impact of the mode change in the Route Reflector (RR) or ASBR context. However, in the PE node role, the change is applied automatically. For RR/ASBR node roles, a manual trigger is necessary.

Command Syntax

```
mpls label mode (vpnv4|vpnv6|all-afs) (all-vrfs|vrf WORD) (per-prefix|per-vrf)
no mpls label mode (vpnv4|vpnv6|all-afs) (all-vrfs|vrf WORD) (per-prefix)

mpls label mode 6pe per-prefix
no mpls label mode 6pe per-prefix
```

Parameters

vrf WORD

Enter a string to identify the VRF

all-vrfs

All the VRFs

per-prefix

Unique MPLS labels are allocated per VPN route in a VRF

per-vrf

Single MPLS labels are allocated for all VPN routes in a VRF

all-afs

All the address families

Default

Enabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 3.0.

Example

```
#configure terminal
(config)#mpls label mode all-afs all-vrfs per-vrf

(config)#no mpls label mode all-afs all-vrfs

(config)#mpls label mode 6pe per-prefix

(config)#no mpls label mode 6pe per-prefix
```

mpls local-packet-handling

Use this command to enable the labeling of locally generated TCP packets. All other locally generated packets are not looked at by the MPLS Forwarder

Use the `no` parameter with this command to disable labeling of locally generated TCP packets.

Command Syntax

```
mpls local-packet-handling
no mpls local-packet-handling
```

Default

By default, mpls local packet handling is disabled

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#mpls local-packet-handling
```

mpls lsp-encap-dscp-preserve

Use this command to preserve DSCP for IP packets encapsulated into MPLS headers when dscp is remarked on access interface. By default, DSCP is not preserved for IP packets encapsulated into MPLS headers.

Use the no parameter with this to unconfigure DSCP preserve.

Command Syntax

```
mpls lsp-encap-dscp-preserve  
no mpls lsp-encap-dscp-preserve
```

Default

Disabled

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 6.5.2.

Examples

```
#configure terminal  
(config)#mpls lsp-encap-dscp-preserve  
(config)#commit  
#configure terminal  
(config)#no mpls lsp-encap-dscp-preserve  
(config)#commit
```

mpls lsp-model uniform

Use this command to configure the MPLS LSP model as uniform.

Use the `no` parameter with this command to configure the MPLS LSP model as pipe or short-pipe.

Command Syntax

```
mpls lsp-model uniform
no mpls lsp-model uniform
```

Parameter

None

Default

By default, model configuration is pipe for XGS devices.

Qumran devices have the following default behavior:

- For L3VPN services, model is short-pipe by default and pipe model can be set by configuring policy-maps with match exp and set queue.
- For L2VPN services, short-pipe model is not supported and the default model is pipe.
- For L2VPN services with rewrite action pop, cos value will always be updated from qos profile irrespective of model.
- For L2VPN services without rewrite, uniform model command does not take effect until 'qos remark-cos' is configured on egress interface.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#mpls lsp-model uniform
(config)#exit

#configure terminal
(config)#no mpls lsp-model uniform
(config)#exit
```

mpls lsp-stitching

Use this command to stitch the LSP segment for an FEC created via a different label signaling protocol.

Use the `no` form of this command to disable this configuration.

Command Syntax

```
mpls lsp-stitching
no mpls lsp-stitching
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 3.0.

Command Example

```
#configure terminal
(config)#mpls lsp-stitching
```

mpls map-route

Use this command to map a prefix to an FEC.

Use the `no` parameter with this command to disable this configuration.

Command Syntax

```
mpls map-route (A.B.C.D/M|A.B.C.D A.B.C.D) (A.B.C.D/M|A.B.C.D A.B.C.D)
no mpls map-route (A.B.C.D/M|A.B.C.D A.B.C.D)
```

Parameters

A.B.C.D

IPv4 prefix to map

A.B.C.D/M

IPv4 prefix to map, plus mask

A.B.C.D

Mask for IPv4 prefix to map

A.B.C.D/M

Mask for IPv4 prefix to map, plus mask.

A.B.C.D

IPv4 forwarding equivalence class for route to map

A.B.C.D

Mask for IPv4 forwarding equivalence class

Default

Disabled

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

In the following examples 5.6.7.8/32 is the FEC for an LSP, and 1.2.3.4 is the prefix to be mapped.

```
#configure terminal
(config)#mpls map-route 1.2.3.4/32 5.6.7.8/32

#configure terminal
(config)#mpls map-route 1.2.3.4 255.255.255.255 5.6.7.8 255.255.255.255
```

mpls min-label-value

Use this command to configure minimum and maximum label value for a label space. Use module names (rsvp | ldp | bgp) to configure minimum and maximum label value for module in a label space, minimum and maximum label space value for a module should be within the range of label space being used. After setting minimum and maximum label value for a label space, make sure to bind the label space to an interface.

Use the `no` parameter with this command to use the default minimum and maximum label value for all the label pools.



Notes:

- The system allows label-space range (maximum and minimum label values) changes for interface-specific label spaces only. The platform-wide label-space range cannot be modified.
- Only label-space 0 (global) is supported. Any label-space other than 0, is not supported.

Command Syntax

```
mpls (rsvp|ldp|bgp) min-label-value <16-1048575> max-label-value <16-1048575> (label-space <0-60000>|)  
no mpls min-label-value max-label-value (label-space <0-60000>|)  
no mpls (rsvp|ldp|bgp) (label-space <0-60000>|)
```

Parameters

rsvp

Label range value for RSVP

ldp

Label range value for LDP

bgp

Label range value for BGP

min-label-value

Specify the minimum label value

<16-1048575>

Minimum size to be used for label pools or protocol range

max-label-value

Specify the maximum label value

<16-1048575>

Maximum size for all label pools

label-space

Label space for which the minimum value needs to be modified

<0-60000>

Range for label space

Default

16

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#mpls min-label-value 50000 max-label-value 80000 label-space 0
```

mpls propagate-ttl

Use this command to enable TTL propagation. Enabling TTL propagation causes the TTL value in the IP header to be copied onto the TTL field in the shim header, at the LSP ingress.

Use the `no` parameter with this command to disable TTL propagation.

Command Syntax

```
mpls propagate-ttl  
no mpls propagate-ttl
```

Parameters

None

Default

Disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#mpls propagate-ttl  
  
#configure terminal  
(config)#no mpls propagate-ttl
```

mpls traffic-eng srlg

Use this command to create a Shared Risk Link Group (SRLG). An SRLG uses secondary backup LSPs or Fast Reroute bypass/detour LSPs that minimize the probability of "fate sharing" with the path of the primary LSP.

Use the `no` form of this command to remove an SRLG.



Notes:

- An interface can be part of multiple SRLG groups upto a maximum of 255 SRLG groups.
- Any addition or deletion of SRLG value on an interface will not recalculate Primary/Backup. It is advised to configure SRLG values before bringing UP RSVP sessions or clear rsvp sessions after updating SRLG values.
- IS-IS supports a maximum of eight risk-groups; configuring more than eight risk-groups is not recommended.
- OSPF supports up to 255 risk-groups.

Command Syntax:

```
mpls traffic-eng srlg <0-4294967295>  
no mpls traffic-eng srlg <0-4294967295>
```

Parameters

<0-4294967295>

Risk group number

Command Mode

Interface mode

Example

```
#configure terminal  
(config)#int eth1  
(config-if)#mpls traffic-eng srlg 1
```

pce entity

Use this command to:

- Assign a PCE entity to a segment routing policy
- Assign a PCE entity to a RSVP tunnel

Use the `no` form of this command to:

- Remove a PCE entity from a segment routing policy
- Remove a PCE entity from a RSVP tunnel

Syntax

```
pce entity <1-255>
no pce entity
```

Parameters

entity <1-255>

Path computation element entity identifier.

Command Mode

RSVP trunk mode

Segment routing policy mode

Applicability

Introduced in OcNOS version 4.0.

Example

```
#configure terminal
(config)#rsvp-trunk T1 ipv4
(config-trunk)#pce entity 127

#configure terminal
(config)#segment-routing
(config-sr)#traffic-engineering
(config-sr-te)#policy P1
(config-sr-pol)#pce entity 127
```

pce lsp-delegate

Use this command to delegate all LSPs of a segment routing policy or RSVP trunk to a PCE entity.

Use `no` form of this command to revoke all LSPs of a segment routing policy or RSVP trunk from a PCE entity.



Note: In case of segment-routing policy, when delegation is revoked, OcNOS (PCC) continues to use the current Explicit Route Object (ERO) if it is available. If the ERO is unavailable, the LSP state remains down. In both cases, the system will simultaneously initiate a new Path Computation Request (PCReq) packet to the Path Computation Element (PCE) to obtain an updated path.

Syntax

```
pce lsp-delegate
no pce lsp-delegate
```

Parameters

None

Default

LSPs are not delegated to a segment routing policy of RSVP trunk.

Command Mode

RSVP trunk mode

Segment routing policy mode

Applicability

Introduced in OcNOS version 4.0.

Example

```
#configure terminal
(config)#segment-routing
(config-sr)#traffic-engineering
(config-sr-te)#policy P1
(config-sr-pol)#pce lsp-delegate

#configure terminal
(config)#rsvp-trunk T1 ipv4
(config-trunk)#pce lsp-delegate
```

pce state-report

Use this command to enable sending a PCEP state report.

Use the `no` form of this command to disable sending a PCEP state report.

Syntax

```
pce state-report
no pce state-report
```

Parameters

None

Default

Sending a PCEP state report is disabled.

Command Mode

RSVP trunk mode

Segment routing policy mode

Applicability

Introduced in OcNOS version 4.0.

Example

```
#configure terminal
(config)#segment-routing
(config-sr)#traffic-engineering
(config-sr-te)#policy P1
(config-sr-pol)#pce state-report

#configure terminal
(config)#rsvp-trunk T1 ipv4
(config-trunk)#pce state-report
```


ping mpls

Use this command to start sending MPLS request packets using various parameters as defined below. Ping packets can be configured for LDP, RSVP, L2 circuit, VPLS, L3 VPN, Segment Routing or generic FEC types.

Command Syntax

```
ping mpls (ldp A.B.C.D/M|rsvp (tunnel-name NAME|egress A.B.C.D)|l2-circuit <1-4294967295> peer
A.B.C.D/M|vpls <1-10000> peer A.B.C.D/M|l3vpn VRFNAME A.B.C.D/M|ipv4 A.B.C.D/M) ({reply-mode
(1|2)|flags|destination A.B.C.D|source A.B.C.D|ttl <1-255>|timeout <1-500>|repeat <5-5000>|interval
<2-20000>|force-explicit-null|detail|})
ping mpls (l3vpn (VRFNAME A.B.C.D/M X:X::X:X/M source A.B.C.D destination A.B.C.D)) ({timeout <1-
500>|ttl <1-255>|repeat <5-5000>|interval <2-20000>|detail|})
ping mpls (6pe default X:X::X:X/M source A.B.C.D destination A.B.C.D)) ({timeout <1-500>|ttl <1-
255>|repeat <5-5000>|interval <2-20000>|detail|})
ping mpls (| protocol-origin (local | pcep) ) (sr-policy SR_POLICY_NAME (| candidate-path CANDIDATE_
PATH_ID)) ({flags|source A.B.C.D | ttl <1-255> | timeout <1-500> | repeat <5-5000>|interval <2-20000>
| force-explicit-null | detail|})
ping mpls (ospf-sr | isis-sr) (ipv4 A.B.C.D/M) ({reply-mode (1|2|3)|flags | destination
A.B.C.D|source A.B.C.D | ttl <1-255> | timeout <1-500> | repeat <5-5000>|interval <2-20000> | force-
explicit-null | detail|})
```

Parameters

ldp

FEC type is LDP

A.B.C.D/M

LDP prefix address

rsvp

FEC type is RSVP

tunnel-name

RSVP tunnel name

NAME

Tunnel name string

egress

RSVP tunnel egress

A.B.C.D

3RSVP tunnel egress address

l2-circuit

FEC type is L2 circuit

<1-4294967295>

L2 circuit ID

peer A.B.C.D/M

IP address of the peer

vpls

FEC type is MPLS VPLS (L2-VPN)

<1-10000>

VPLS instance ID

peer

VPLS peer

A.B.C.D/M

VPLS peer address

l3vpn

FEC type is MPLS VPN (L3-VPN)

VRFNAME

VPN instance name

A.B.C.D./M

VPN prefix

X:X::X:X/M

VPNv6 prefix

6pe

FEC type (6PE)

default

VPN Instance Name (default)

X:X::X:X/M

6PE Prefix

ipv4

FEC type is generic; use for static/SNMP label switched paths

A.B.C.D/M

IPv4 prefix address

Protocol-origin

Protocol origin for SR policy

Local

local originated sr policy

Pcep

Pcep originated SR policy

SR policy

SR policy ping

SR_POLICY_NAME

Policy name

Candidate-path

Candidate path

CANDIDATE_PATH_ID

Candidate path id

reply-mode

Reply mode, one of

1

Do not reply

2

Reply via UDP/IP packet (default)

flags

Validate FEC stack

destination

Destination address

A.B.C.D

IPv4 address of the destination

source

Source address

A.B.C.D

IPv4 address of the source

ttl

Trace packet Time-to-live

<1-255>

Trace packet TTL value

repeat

Repeat sending of ping packets

<5-5000>

Number of pings to send

interval

Interval between ping packets, in milliseconds

<2-20000>

Interval value

timeout

Time to wait before rejecting the probe as a failure, in seconds

<1-500>

Timeout value

force-explicit-null

Force Explicit NULL label

detail

Print detailed output of the ping

Defaults

TTL value is 255.

Timeout value is 60 seconds.

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3. The SR ping command was introduced in OcNOS version 4.1.

Example

```
#ping mpls ipv4 10.10.0.0/24 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 ttl 226
timeout 65 repeat 6 interval 3 detail force-explicit-null

#ping mpls l2-circuit 3 peer 192.0.2.0/32 reply-mode 2 flags destination 127.1.3.4 source 10.10.0.1
ttl 226 timeout 65 repeat 6 interval 3 detail force-explicit-null

#ping mpls l3vpn vrfa 10.10.0.0/24 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 ttl 226
```

```
timeout 65 repeat 6 interval 3 detail force-explicit-null

#ping mpls ldp 10.10.0.0/24 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 ttl 226 timeout
65 repeat 6 interval 3 detail force-explicit-null

#ping mpls rsvp egress 1.2.3.5 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 ttl 226
timeout 65 repeat 6 interval 3 detail force-explicit-null

#ping mpls rsvp tunnel-name tun1 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 ttl 226
timeout 65 repeat 6 interval 3 detail force-explicit-null

#ping mpls vpls 2 peer 10.10.0.0 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 ttl 226
timeout 65 repeat 6 interval 3 detail force-explicit-null

# ping mpls sr-policy LSP3_R1_R4_R2 detail
# ping mpls ospf-sr ipv4 2.2.2.2/32 detail
# ping mpls isis-sr ipv4 2.2.2.2/32 detail

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

! seq_num = 1 200.0.0.1 2.02 ms
! seq_num = 2 200.0.0.1 2.00 ms
! seq_num = 3 200.0.0.1 1.93 ms
! seq_num = 4 200.0.0.1 2.14 ms
! seq_num = 5 200.0.0.1 1.78 ms

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 1.78/1.96/2.14
```

rewrite ingress

Use this command to configure a match VLAN action for a service template.

Use the `no` parameter with this command to remove a match VLAN action for a service template.

Command Syntax

```
rewrite ingress (((pop |translate <2-4094>)(|outgoing-tpid (dot1.ad |dot1.q)))) (push <2-4094>))
no rewrite ingress (pop |push |translate)
```

Parameters

pop

POP the outer VLAN identifier from ACCESS->NETWORK and PUSH the match outer VID to NETWORK->ACCESS

translate

Translate the outer VLAN identifier to configured action VID for ACCESS->NETWORK and translate to the match outer VID for NETWORK->ACCESS

<2-4094>

Outer VLAN identifier

outgoing-tpid

Outgoing TPID, set the outer-tpid for the NETWORK->ACCESS

dot1.ad

Set TPID value as 0x88a8 for the traffic NETWORK->ACCESS

dot1.q

Set TPID value as 0x8100 for the traffic NETWORK->ACCESS

push

PUSH the outer VLAN identifier from ACCESS->NETWORK and POP the Outer VID from NETWORK->ACCESS

<2-4094>

Outer VLAN identifier

Command Mode

MPLS SVC mode

Applicability

This command was introduced in OcNOS version 1.3.4, and changed in OcNOS version 3.0.

Examples

```
#configure terminal
(config)#service-template C2
(config-svc)#match double-tag outer-vlan 9 inner-vlan 8
(config-svc)#rewrite ingress translate 7 outgoing-tpid dot1.ad
(config-svc)#exit

(config)#service-template C2
(config-svc)#no rewrite ingress translate
(config-svc)#exit
```

secondary-priority srlg-disjoint

Use this command to set how to avoid the SRLGs (Shared Risk Link Groups) of a protected primary.

A fast-reroute/secondary path for an LSP that is disjoint from the primary ensures that a single point of failure on a particular link does not bring down both the primary and secondary paths in the LSP.



Note: The SRLG option configured in RSVP-TRUNK mode (this command) takes higher preference than the option configured in RSVP router mode (see the [srlg-disjoint \(page 1483\)](#) command).

Use the `no` form of this command to not avoid the SRLGs of a protected interface.

Command Syntax

```
secondary-priority <1-5> srlg-disjoint (forced|preferred)
no secondary-priority <1-5> srlg-disjoint
```

Parameters

forced

The router does not create the secondary/backup tunnel unless it avoids SRLGs of the primary-path/protected-interface.

preferred

With two explicit paths, the first explicit path tries to avoid the SRLGs of the primary-path/ protected interface. If that does not work, the secondary/backup tunnel uses the second path (which ignores SRLGs).

Command Mode

RSVP -TRUNK mode

Example

```
#configure terminal
(config)#rsvp-trunk t1
(config-rsvp)# secondary-priority 1 srlg-disjoint forced
```

secondary srlg-disjoint

Use this command to set how to avoid the SRLGs (Shared Risk Link Groups) of a protected primary.

A fast-reroute/secondary path for an LSP that is disjoint from the primary ensures that a single point of failure on a particular link does not bring down both the primary and secondary paths in the LSP.



Note: The SRLG option configured in RSVP-TRUNK mode (this command) takes higher preference than the option configured in RSVP router mode (see 'srlg-disjoint').

Use the `no` form of this command to not avoid the SRLGs of a protected interface.

Command Syntax

```
secondary srlg-disjoint (forced|preferred)
no secondary srlg-disjoint
```

Parameters

forced

The router does not create the secondary/backup tunnel unless it avoids SRLGs of the primary-path/protected-interface.

preferred

With two explicit paths, the first explicit path tries to avoid the SRLGs of the primary-path/ protected interface. If that does not work, the secondary/backup tunne uses the second path (which ignores SRLGs).

Command Mode

RSVP -TRUNK mode

Example

```
#configure terminal
(config)#rsvp-trunk t1
(config-rsvp)# secondary srlg-disjoint forced
```

service-tpid

Use this command to configure service tpid for the MPLS layer-2 virtual circuit.

Use the no parameter with this command to delete service tpid from the MPLS layer-2 virtual circuit.

Command Syntax

```
service-tpid {dot1.q|dot1.ad|0x9100}  
no service-tpid
```

Parameters

0x9100

Set tpid value as 0x9100

dot1.ad

Set tpid value as 0x88a8

dot1.q

Set tpid value as 0x8100

Default

Disabled

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 1.3

Example

```
#configure terminal  
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4  
(config-pseudowire)#service-tpid dot1.ad
```


show admin-groups

Use this command to display all configured administrative groups.

Command Syntax

```
show admin-groups
```

Parameters

None

Command Mode

Execution mode and Privilege mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following sample shows the output of the `show admin-group` command.

```
#show admin-groups
Admin group detail:
Value of 0 associated with admin group 'a'
Value of 1 associated with admin group 'b'
Value of 2 associated with admin group 'c'
Value of 4 associated with admin group 'd'
```

The following table explains the show command output fields.

Table 54. show admin-groups output field

| Field | Description |
|--------------------|---|
| Admin group detail | Display information about configured Multi Protocol Label Switching (MPLS) administrative groups. |

show hsl mpls l3vpn-ftn

Use this command to view the L3VPN FTN details in tabular form. It includes information on all the multipath FTNs.

Command Syntax

```
show hsl mpls l3vpn-ftn
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3. Enhanced the command for ECMP in OcNOS version 6.6.1

Examples

```
#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| VRF |      DESTINATION      | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref |Prefix|
| MPLS |      OUT      | OUT      | NEXTHOP      |         |     |      |
| ID   |      FEC        | cnt | cnt | Ifname | IFNAME | LABEL |
| OBJECT | cnt | cnt |         |         |         |
+-----+-----+-----+-----+-----+-----+-----+-----+
2      B> 192.12.4.0/24      - /-                                0x20001011 00:00:07      1      1
/      /10
tp500006 xe27      27520      22.22.22.22      0x20001011 00:00:07      1      1
mpls-
```

show mpls bandwidth-class

Use this command to view bandwidth class parameters: bandwidth class name; allocated bandwidth; setup hold priority

Command Syntax

```
show mpls bandwidth-class
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
show mpls bandwidth-class
Bandwidth-class: BW_1
Bandwidth: 6k          Setup-priority: 1  Class-type: 1
```

The following table explains the show command output fields.

Table 55. show mpls bandwidth-class output field

| Field | Description |
|-----------------|--|
| Bandwidth-class | Bandwidth for each class type. |
| Bandwidth | Bandwidth configured for the active MPLS. |
| Setup-Priority | The setup priority is compared with other setup priorities for established sessions on the link to determine whether some of them should be preempted to accommodate the new session. Sessions with lower hold priorities are preempted. |
| Class-type | Bandwidth allocated for the specified class type. |

show mpls counters ldp

Use this command to display traffic statistics for FTNs and ILMs configured by LDP.

Command Syntax

```
show mpls counters ldp ((ftn (|A.B.C.D/M)) | (ilm (|A.B.C.D/M)) |)
```

Parameter

- ftn**
FEC-to-NHLFE map counters
- A.B.C.D/M**
FEC prefix
- ilm**
Incoming label map counters
- A.B.C.D/M**
FEC prefix

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 1.3.1.



Note: For Qumran, counters are not available for transit nodes.

Examples

```
OcNOS#show mpls counter ldp
[FTN statistics]
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
|      FEC      | Pri | nhlfe_ix | out-label | out-intf |   nexthop   |      Tx      |
| packets  | Tx bytes  |             |             |             |             |             |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 4.4.4.4/32    | 1    | 1        | 3         | xe3       | 10.10.1.2   | 62          | 4
484
VPN AGGR
STAT  NA    NA    NA    NA    NA    64500    9417000

[ILM statistics]
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
|      FEC      | in-label | out-label | Rx packets | Rx bytes |      Tx      |
| packets  | Tx bytes  |             |             |             |             |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
```

The following table explains the show command output fields.

Table 56. show mpls counters ldp output field

| Field | Description |
|----------------|---|
| FTN statistics | Displays the statistics details of FTN. |
| ILM statistics | Displays the statistics details of ILM. |
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| In-label | Displays the ingress (incoming interface) label for this segment. |
| Out-label | Displays the egress (outgoing interface) label for this segment. |
| Rx packets | Number of hello packets received from the neighbor. |
| Rx bytes | Size of hello packets received from the neighbor. |
| Tx packets | Number of hello packets sent to the neighbor. |
| Tx bytes | Size of hello packets sent to the neighbor. |

show mpls

Use this command to display MPLS data.

Command Syntax

```
show mpls
```

Parameters

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3. This command has been enhanced to include MPLS LSP models in OcNOS version 7.0.0.

Examples

The following subsection displays a variety of `show mpls` commands.

```
#show mpls
Minimum label configured: 16
Maximum label configured: 1048575
Per label-space information:
  Label-space 0 is using minimum label: 16 and maximum label: 1048575
Custom ingress TTL configured: 64
Custom egress TTL configured: none
Log message detail: none
Admin group detail: none
MPLS LSP Model
  transport : short-pipe
  l2vpn     : pipe
  l3vpn     : short-pipe
#
```

The following table explains the show command output fields.

Table 57. show mpls output field

| Field | Description |
|--------------------|--|
| Packets dropped IP | Displays the number of packets dropped over the internet protocol. |
| Dropped MPLS | Displays the number of packets dropped over the MPLS. |
| Sent to IP | Displays the number of packets transmitted to the internet protocol. |
| Labeled | Number of labeled packets in the interface. The MPLS-labeled packets are switched after a label lookup/switch instead of a lookup into the IP table. Labels of pop-and-forward mpls tunnel: P—Pop labels. |

Table 57. show mpls output field (continued)

| Field | Description |
|----------|--|
| | D—Delegation labels. |
| Switch | Type of switching on the links needed for the MPLS. |
| Class | Creates a class map to be used for matching traffic to a specified class, and enters class-map configuration mode. |
| MPLS LSP | Label Switched Path model used to transport data. |

show mpls counters static

Use this command to display traffic statistics for statically configured FTNs and ILMs.

Command Syntax

```
show mpls counters static ((ftn (A.B.C.D/M|)) | (ilm (A.B.C.D/M|)) |)
```

Parameter

ftn
FEC-to-NHLFE map counters
A.B.C.D/M
FEC prefix

ilm
Incoming label map counters
A.B.C.D/M
FEC prefix

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 1.3.1.



Note: For Qumran, counters are not available for transit nodes.

Examples

```
#show mpls counters static
[ FTN statistics ]
+-----+-----+-----+-----+
|      FEC      | out-label | Tx packets | Tx bytes |
+-----+-----+-----+-----+
| 192.168.1.0/24 | 100       | 0          | 0        |
| 192.168.2.0/24 | 200       | 0          | 0        |
+-----+-----+-----+-----+

[ ILM statistics ]
+-----+-----+-----+-----+-----+-----+
|      FEC      | in-label | out-label | Rx packets | Rx bytes | Tx
| packets      | Tx bytes |           |            |          |
+-----+-----+-----+-----+-----+-----+
| 0.0.0.0/0     | 201      | n/a      | 0          | 0        | n/a
| n/a          |          |          |            |          |
| 0.0.0.0/0     | 101      | n/a      | 0          | 0        | n/a
| n/a          |          |          |            |          |
| 192.168.3.0/24 | 301      | 302      | 0          | 0        | 0
| 0            |          |          |            |          |
| 192.168.4.0/24 | 401      | 402      | 0          | 0        | 0
```



```
0
#
```

The following table explains the show command output fields.

Table 58. show mpls counters static output field

| Field | Description |
|----------------|---|
| FTN statistics | Displays the statistics details of FTN. |
| ILM statistics | Displays the statistics details of ILM. |
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| In-label | Displays the ingress (incoming interface) label for this segment. |
| Out-label | Displays the egress (outgoing interface) label for this segment. |
| Rx packets | Number of hello packets received from the neighbor. |
| Rx bytes | Size of hello packets received from the neighbor. |
| Tx packets | Number of hello packets sent to the neighbor. |
| Tx bytes | Size of hello packets sent to the neighbor. |

show mpls counters rsvp

Use this command to display traffic statistics for LSPs configured by RSVP.

Command Syntax

```
show mpls counters rsvp ((tunnel-name NAME) | (tunnel-id TUNNEL_ID) | (node-role (ingress | transit | egress)) |)
```

Parameter

NAME

RSVP tunnel name

TUNNEL_ID

RSVP tunnel identifier

ingress

LSP role is ingress

transit

LSP role is transit

egress

LSP role is egress

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 1.3.1.



Note: For Qumran, counters are not available for transit nodes.

Examples

```
#show mpls counters rsvp
Tunnel-id 5001 Extended Tunnel-ID 9.9.9.1 Egress 9.9.9.2
  lsp-name : t1-Primary                                [Ingress]
  lsp-ingress : 9.9.9.1                                lsp-id : 101
  Rx pkts : 0                                           Rx bytes : 0
  Tx pkts : 0                                           Tx bytes : 0

  lsp-name : t1-Secondary                                [Ingress]
  lsp-ingress : 9.9.9.1                                lsp-id : 102
  Rx pkts : 0                                           Rx bytes : 0
  Tx pkts : 0                                           Tx bytes : 0

Tunnel-id 5002 Extended Tunnel-ID 9.9.9.1 Egress 9.9.9.3
  lsp-name : t2-Primary                                [Ingress]
  lsp-ingress : 9.9.9.1                                lsp-id : 104
  Rx pkts : 0                                           Rx bytes : 0
  Tx pkts : 0                                           Tx bytes : 0
```

```

lsp-name : t2-Detour                                [Ingress]
lsp-ingress : 1.1.49.1                             lsp-id : 104
Rx pkts : 0                                         Rx bytes : 0
Tx pkts : 0                                         Tx bytes : 0

```

The following table explains the show command output fields.

Table 59. show mpls counters rsvp output field

| Field | Description |
|--------------------|---|
| Tunnel-id | Tunnel identifier (destination port) for the RSVP session. |
| Extended Tunnel-ID | Extended Tunnel identifier (destination port) for the RSVP session. |
| Egress | Egress router is the final MPLS device that removes the last label before packets leave the MPLS network. |
| lsp-name | Name of the SPRING-TE LSP. |
| Ingress | The router at the beginning of an LSP. This router encapsulates IP packets with an MPLS Layer 2 frame and forwards it to the next router in the path. |
| lsp-ingress | The router at the beginning of an LSP. |
| lsp-id | Specify the generic LSP identifier. |
| Rx packets | Number of hello packets received from the neighbor. |
| Rx bytes | Size of hello packets received from the neighbor. |
| Tx packets | Number of hello packets sent to the neighbor. |
| Tx bytes | Size of hello packets sent to the neighbor. |

show mpls cross-connect-table

Use this command to display detailed information for all entries created in the MPLS cross-connect table.

Command Syntax

```
show mpls cross-connect-table
```

Parameters

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

The following is a sample output of the show mpls cross-connect-table

```
#show mpls cross-connect-table
Cross connect ix: 3, in intf: -, in label: 0, out-segment ix: 3
  Owner: RSVP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 3, owner: RSVP, out intf: eth1, out label: 16
  Nexthop addr: 10.10.20.80, cross connect ix: 3, op code: Push

Cross connect ix: 6, in intf: -, in label: 0, out-segment ix: 6
  Owner: RSVP, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 6, owner: RSVP, out intf: eth1, out label: 17
  Nexthop addr: 10.10.20.80, cross connect ix: 6, op code: Push
#
```

The following table explains the show command output fields.

Table 60. show mpls cross-connect-table output field

| Field | Description |
|------------------|---|
| Cross connect ix | Displays the table index for the cross-connect. |
| In intf | Installed as a result of configuring an interface. |
| In label | Displays the ingress (incoming interface) label for this segment. |
| Out-segment ix | Displays the outbound segment index. |
| Owner | Displays the creator of this segment, typically a protocol such as BGP. |
| Persistent | Displays whether the tunnel is persistent – Yes or No. |
| Admin Status | Indicates whether the user can administratively disable a peer while still preserving its configuration. Up = Yes, Down = No. |

Table 60. show mpls cross-connect-table output field (continued)

| Field | Description |
|--------------|---|
| Oper Status | Displays the current status of the cross-connect segment – Up or Down |
| Nexthop addr | Displays the IP address of the next hop. |
| Op code | PUSH = Replace the top label with another and then push one or more additional labels onto the label stack SET = Set the next hop label. |

show mpls cross-connect-table backup

Use this command to display detailed information for all entries created in the MPLS backup cross-connect table.

Command Syntax

```
show mpls cross-connect-table backup
```

Parameter

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command is newly introduced in OcNOS version 5.1.

Example

The following is a sample output of the show mpls cross-connect-table backup

```
#show mpls cross-connect-table backup

Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 17
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 17, owner: LDP, Stale: NO, out intf: xe3, out label: 24324
Nexthop addr: 20.20.20.2 cross connect ix: 2, op code: Push
Primary xc-ix 2 out-segment ix 17

Backup Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 12
Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Not present
Out-segment with ix: 12, owner: LDP, Stale: NO, out intf: xe21, out label: 24320
Nexthop addr: 10.10.10.2 cross connect ix: 2, op code: Push
Primary xc-ix 2 out-segment ix 17
```

show mpls forwarding-table detail

Use this command to view forwarding table entries.

Command Syntax

```
show mpls forwarding-table ((A.B.C.D/M|X:X::X:X/M| algorithm) | count | detail |)
```

Parameters

A.B.C.D/M

FEC IPv4

X:X::X:X/M

FEC IPv6

algorithm

Specifies the Flexible Algorithm

count

Count of IPv4 FTNs.

detail

show detail

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3 and the detail option is applicable form OcNOS version 4.2. Added parameter **algorithm** in OcNOS version 6.6.0.

Detail applicable form is modified to display backup information if present (only for LDP FTNs).

Example

If LFA flag is not enabled

```
#show mpls forwarding-table detail
FEC prefix: 41.41.41.41/32, FTN-ID: 5
  Owner: LDP, FTN type: Regular, State: Installed
  Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
  LSP-ID: N/A, LSP-type: Primary
  NHLFE-id: 7
  Out-Label: 3, Out-Intf: xe1, Nexthop: 50.50.50.1
  Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
  ELC: No
```

If LFA is enabled and backup not present:

```
FEC prefix: 41.41.41.41/32, FTN-ID: 5
  Owner: LDP, FTN type: Regular, State: Installed
  Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
  LSP-ID: N/A, LSP-type: Primary
```

```
NHLFE-id: 7
Primary : Out-Label: 3, Out-Intf: xe1, Nexthop: 50.50.50.1
Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
ELC: No
```

If Backup is present:

```
#show mpls forwarding-table detail
FEC prefix: 10.10.10.0/24, FTN-ID: 2
  Owner: LDP, FTN type: Regular, State: Installed
  Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
  LSP-ID: N/A, LSP-type: Primary
  NHLFE-id: 4
  Primary : Out-Label: 3, Out-Intf: xe1, Nexthop: 50.50.50.1
  LFA Backup : Out-Label: 3, Out-Intf: xe3, Nexthop: 20.20.20.1
  Primary : Out-Label: 3, Out-Intf: xe3, Nexthop: 20.20.20.1
  LFA Backup : Out-Label: 3, Out-Intf: xe1, Nexthop: 50.50.50.1
  Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
  ELC: No
```

If algo is configured

```
show mpls forwarding-table detail algorithm 128
FEC prefix: 25.0.1.2/32, FTN-ID: 7
  Owner: ISIS-SR-FA, FTN type: REGULAR, State: Installed
  Tunnel-Name: N/A, Tunnel-id: N/A, Color: N/A
  LSP-ID: N/A, LSP-type: Primary
  ECMP nhlfe-ix: 80
  Primary : Out-Label: 3, Out-Intf: xe5, Nexthop: 11.0.12.2
  TI-LFA Backup : Out-Label: 21282, Out-Intf: xe22, Nexthop: 11.0.15.5
  Exp-bits: 0x0, Incoming DSCP: none, QoS Resource id: 0
  ELC: Yes
```


show mpls forwarding-table

Use this command to view forwarding table entries.

Command Syntax

```
show mpls forwarding-table ((A.B.C.D/M|X:X::X:X/M|algorithm) | count | detail |)
```

Parameters

A.B.C.D/M

FEC IPv4

X:X::X:X/M

FEC IPv6

algorithm

Specifies the Flexible Algorithm

count

Count of IPv4 FTNs.

detail

Show detail

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3 and the `detail` parameter added in OcNOS version 4.2. Added parameter **algorithm** in OcNOS version 6.6.0 The output was extended to display backup Information as well.

Example

If Backup not present:

```
#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out-Intf | ELC | Nexthop |
|---------|----------------|--------|----------|------------|-----|----------|-----------|----------|-----|---------|
| L> | 10.10.10.0/24 | 2 | 4 | - | Yes | LSP_ | | | | |
| DEFAULT | 3 | xe1 | No | 50.50.50.1 | | | | | | |
| | | | | - | Yes | LSP_ | | | | |
| DEFAULT | 3 | xe3 | No | 20.20.20.1 | | | | | | |
| L> | 30.30.30.0/23 | | 3 | 5 | - | Yes | LSP_ | | | |
| DEFAULT | 3 | xe1 | No | 50.50.50.1 | | | | | | |
| L> | 30.30.30.0/24 | | 1 | 2 | | Yes | LSP_ | | | |
| DEFAULT | 3 | xe6 | No | 40.40.40.2 | | | | | | |
| L> | 35.35.35.35/32 | | 5 | 7 | | Yes | LSP_ | | | |
| DEFAULT | 3 | xe6 | No | 40.40.40.2 | | | | | | |
| L> | 41.41.41.41/32 | | 4 | 6 | - | Yes | LSP_ | | | |

```

DEFAULT 3          xe1      No    50.50.50.1
L> 44.44.44.44/32    6      11    -          Yes  LSP_
DEFAULT 3          xe3      No    20.20.20.1
#

```

If Backup is present:

```

#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
        B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport

Code  FEC              FTN-ID  Nhlfe-ID  Tunnel-id  Pri  LSP-Type  Out-Label  Out-
Intf  ELC  Nexthop
L> 10.10.10.0/24      2      4      -          Yes  LSP_
DEFAULT 3          xe1      No    50.50.50.1
                        -          No    LSP_
DEFAULT 3          xe3      No    20.20.20.1
                        -          Yes  LSP_
DEFAULT 3          xe3      No    20.20.20.1
                        -          No    LSP_
DEFAULT 3          xe1      No    50.50.50.1
L> 30.30.30.0/23      3      5      -          Yes  LSP_
DEFAULT 3          xe1      No    50.50.50.1
L> 30.30.30.0/24      1      2      -          Yes  LSP_
DEFAULT 3          xe6      No    40.40.40.2
L> 35.35.35.35/32    5      7      -          Yes  LSP_
DEFAULT 3          xe6      No    40.40.40.2
L> 41.41.41.41/32    4      6      -          Yes  LSP_
DEFAULT 3          xe1      No    50.50.50.1
L> 44.44.44.44/32    6      11    -          Yes  LSP_
DEFAULT 3          xe3      No    20.20.20.1
                        -          No    LSP_
DEFAULT 24326        xe1      No    50.50.50.1
L> 1.1.1.1/32        2      4      -          Yes  LSP_
DEFAULT 24320        xe12     No    20.1.1.1
                        11     -          No    LSP_
DEFAULT 24320        xe5      No    30.1.1.1
                        11     -          Yes  LSP_
DEFAULT 24320        xe5      No    30.1.1.1
                        3      -          No    LSP_
DEFAULT 24320        xe12     No    20.1.1.1
L> 2.2.2.2/32        3      6      -          Yes  LSP_
DEFAULT 3          xe12     No    20.1.1.1
                        15     -          No    LSP_
DEFAULT 24324        xe5      No    1.1.1.1
                                    (via
30.1.1.1 ,label 24320)
L> 3.3.3.3/32        1      2      -          Yes  LSP_
DEFAULT 3          xe5      No    30.1.1.1
                        16     -          No    LSP_
DEFAULT 24325        xe12     No    1.1.1.1
                                    (via
20.1.1.1 ,label 24320)

```

If algo is configured.

```

#show mpls forwarding-table algorithm 128
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
        B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

FTN-ECMP LDP: Disabled, SR: Enabled

```

Code  FEC              FTN-ID  Nhlfe-ID  Tunnel-ID  Pri  Out-Label  Out-

```

| Intf | ELC | Nexthop | Algo-Num | UpTime | | | | | |
|-----------|-------------|----------|----------|--------|-----------|-------|------|-----|---|
| i> | 25.0.1.2/32 | 7 | 80 | - | - | - | - | - | - |
| | - | 128 | 20:27:05 | | | | | | |
| | | | 4 | 0 | Yes | 3 | xe5 | Yes | |
| 11.0.12.2 | - | - | 79 | - | | | | | |
| | No | 21282 | xe22 | Yes | 11.0.15.5 | - | - | | |
| i | | | | | | | | | |
| (b)> | 25.0.1.2/32 | 20 | 49 | 2205 | Yes | 21282 | xe5 | No | 1 |
| 1.0.12.2 | 128 | 18:45:40 | | | | | | | |
| i> | 25.0.1.3/32 | 8 | 87 | - | (e) | - | - | - | - |
| | - | 128 | 20:27:05 | | | | | | |
| | | | 26 | 0 | Yes | 21283 | xe5 | Yes | |
| 11.0.12.2 | - | - | 64 | - | | | | | |
| | No | 21283 | xe22 | Yes | 11.0.15.5 | - | - | | |
| | | | 64 | 0 | Yes | 21283 | xe22 | Yes | |
| 11.0.15.5 | - | - | 26 | - | | | | | |
| | No | 21283 | xe5 | Yes | 11.0.12.2 | - | - | | |
| i> | 25.0.1.4/32 | 9 | 89 | - | - | - | - | - | - |
| | - | 128 | 20:27:05 | | | | | | |
| | | | 1 | 0 | Yes | 3 | xe15 | Yes | |
| 11.0.14.4 | - | - | 67 | - | | | | | |
| | No | 21284 | xe22 | Yes | 11.0.15.5 | - | - | | |
| i> | 25.0.1.5/32 | 17 | 92 | - | - | - | - | - | - |
| | - | 128 | 18:45:41 | | | | | | |
| | | | 17 | 0 | Yes | 3 | xe22 | Yes | |
| 11.0.15.5 | - | - | 91 | - | | | | | |
| | No | 21285 | xe5 | Yes | 11.0.12.2 | - | - | | |
| i | | | | | | | | | |
| (b)> | 25.0.1.5/32 | 19 | 44 | 2204 | Yes | 21285 | xe22 | No | 1 |
| 1.0.15.5 | 128 | 18:45:40 | | | | | | | |
| i> | 25.0.1.6/32 | 12 | 98 | - | - | - | - | - | - |
| | - | 128 | 18:45:41 | | | | | | |
| | | | 72 | 0 | Yes | 21286 | xe22 | Yes | |
| 11.0.15.5 | - | - | 86 | - | | | | | |
| | No | 21286 | xe5 | Yes | 11.0.12.2 | - | - | | |

show mpls ftn-table

Use this command to display FTN (FEC-To-NHLF) table information.

Command Syntax

```
show mpls ftn-table((A.B.C.D/M|X:X::X:X/M|algorithm) | out-segment))
```

Parameters

A.B.C.D/M

FEC IPv4

X:X::X:X/M

FEC IPv6

algorithm

Specifies the Flexible Algorithm

out-segment

Specifies multi-segments for the supported protocols (ldp, ospf-sr, isis-sr)

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3. Added parameter **algorithm** in OcNOS version 6.6.0.

Example

If LFA is not present

```
#show mpls ftn-table

Primary FTN entry with FEC: 10.10.10.0/24, id: 2, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: N/A, Stale: NO, out intf: xel, out label: 3
Nexthop addr: 50.50.50.1          cross connect ix: 3, op code: Push
```

If Backup is present

```
#show mpls ftn-table

Primary FTN entry with FEC: 10.10.10.0/24, id: 2, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 3
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: N/A, Stale: NO, out intf: xel, out label: 3
Nexthop addr: 50.50.50.1          cross connect ix: 3, op code: Push
```

```
Backup Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 8
```

```
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 8, owner: N/A, Stale: NO, out intf: xe3, out label: 3
Nexthop addr: 20.20.20.1          cross connect ix: 3, op code: Push
```

If algo is specified

```
show mpls ftn-table 25.0.1.2/32 algorithm 128
Primary FTN entry with FEC: 25.0.1.2/32, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 20:32:18, UpTime: 20:32:18, LastUpdate: 18:50:48
  Owner: ISIS-SR-FA, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none,
  Algorithm Number:128
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 4 refcount: 1
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  State: Active
  Out-segment with ix: 4, owner: N/A, Stale: NO, refcount: 21, out intf: xe5, out label: 3
  Nexthop addr: 11.0.12.2          cross connect ix: 3, op code: Push

  Backup Cross connect ix: 14, in intf: - in label: 0 out-segment ix: 79 bypass ftn-ix: 19
  Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
  State: Active
  Out-segment with ix: 79, owner: ISIS-SR-FA, Stale: NO, refcount: 2, out intf: xe22, out label:
21282
  Nexthop addr: 11.0.15.5          cross connect ix: 14, op code: Push

Primary FTN entry with FEC: 25.0.1.2/32, id: 20, row status: Active, Tunnel-Policy: N/A, State:
Installed
  CreateTime: 18:50:53, UpTime: 18:50:53, LastUpdate: N/A
  Owner: ISIS-SR-FA, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP:
none, Algorithm Number:128
  Tunnel id: 2205, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
  Cross connect ix: 37, in intf: - in label: 0 out-segment ix: 49 refcount: 1
  Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
  State: Active
  Out-segment with ix: 49, owner: ISIS-SR-FA, Stale: NO, refcount: 1, out intf: xe5, out label:
21282
  Nexthop addr: 11.0.12.2          cross connect ix: 37, op code: Push
```

show mpls ilm-table

Use this command to view Incoming label mapping (ILM) table entries.

Command Syntax

```
show mpls ilm-table ((A.B.C.D/M| VRFNAME|X:X::X:X/M|algorithm) (count|))
```

Parameters

A.B.C.D/M

FEC IPv4

VRFNAME

Shows MPLS Incoming label map table by vrf name.

X:X::X:X/M

FEC IPv6

algorithm

Specifies the Flexible Algorithm

count

Count of entries in ILM table.

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy,          U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled
Code  FEC/VRF/L2CKT  ILM-ID  In-Label  Out-Label  In-Intf  Out-
Intf/VRF  Nexthop  pri  Algo-Num  UpTime  UPStr  peers
L>  192.0.1.0/24    16      24335    24334    N/A      xe8.200    22.26.1.26
      Yes N/A      00:42:48  3      24335    24334    N/A      xe8.100    22.26.0.26
      Yes -        -
L>  26.27.0.0/24    8      24327    0        N/A      xe8.200    22.26.1.26
      Yes N/A      00:42:48  3      24327    0        N/A      xe8.100    22.26.0.26
      Yes -        -
L>  9.9.0.100/32    4      24323    24320    N/A      xe8.200    22.26.1.26
      Yes N/A      00:42:48  3      24323    24320    N/A      xe8.100    22.26.0.26
      Yes -        -
L>  9.9.2.100/32    2      24321    Nolabel  N/A      N/A        127.0.0.1
```

```

      Yes  N/A      00:42:49  3
L>  9.9.1.100/32    1      24320      Nolabel      N/A      N/A      127.0.0.1
      Yes  N/A      00:42:49  3
L>  9.9.3.100/32    3      24322      Nolabel      N/A      N/A      127.0.0.1
      Yes  N/A      00:42:49  3
L>  24.26.0.0/24    6      24325      0      N/A      xe8.200      22.26.1.26
      Yes  N/A      00:42:48  3
      24325      0      N/A      xe8.100      22.26.0.26
      Yes  -      -

```

If also is specified

```

show mpls ilm-table algorithm 128
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
      K - CLI ILM, T - MPLS-TP, s - Stitched ILM
      S - SNMP, L - LDP, R - RSVP, C - CRLDP
      B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
      O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
      P - SR Policy,      U - unknown, UPStr - upstream

```

ILM-ECMP LDP: Disabled, SR: Enabled

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out-UPStr | peers |
|----------|---------------|----------|----------|-----------|---------|-----------|-----------|
| Intf/VRF | Nextthop | | pri | Algo-Num | UpTime | | |
| i> | 25.0.1.2/32 | 19 | 17282 | 3 | N/A | xe5 | 11.0.12.2 |
| | Yes 128 | 20:41:58 | 17282 | 21282 | N/A | xe22 | 11.0.15.5 |
| | No - | - | | | | | |
| i> | 25.0.1.1/32 | 6 | 17281 | Nolabel | N/A | N/A | 127.0.0.1 |
| | Yes 128 | 20:00:05 | | | | | |
| i> | 25.0.1.4/32 | 21 | 17284 | 3 | N/A | xe15 | 11.0.14.4 |
| | Yes 128 | 20:41:58 | 17284 | 21284 | N/A | xe22 | 11.0.15.5 |
| | No - | - | | | | | |
| i> | 25.0.1.3/32 | 20 | 17283 | 21283 | N/A | xe5 | 11.0.12.2 |
| | Yes 128 | 20:41:58 | 17283 | 21283 | N/A | xe22 | 11.0.15.5 |
| | No - | - | 17283 | 21283 | N/A | xe22 | 11.0.15.5 |
| | Yes - | - | 17283 | 21283 | N/A | xe5 | 11.0.12.2 |
| | No - | - | | | | | |
| i> | 25.0.1.6/32 | 24 | 17286 | 21286 | N/A | xe22 | 11.0.15.5 |
| | Yes 128 | 20:41:58 | 17286 | 21286 | N/A | xe5 | 11.0.12.2 |
| | No - | - | | | | | |
| i> | 25.0.1.5/32 | 33 | 17285 | 3 | N/A | xe22 | 11.0.15.5 |
| | Yes 128 | 19:00:34 | 17285 | 21285 | N/A | xe5 | 11.0.12.2 |
| | No - | - | | | | | |

The following table shows the status codes displayed at the start of a route entry.

Table 61. status code output field

| Status Code | Field | Description |
|-------------|---------------|---|
| | Installed ILM | Number of ILM entry installed. |
| * | Selected ILM | ILM entry selected in the interface. |
| P | Stale ILM | Stale marked ILM due to on-going graceful restart of MPLS module. |
| K | CLI ILM | Admin configured static ILM entry. |
| T | MPLS-TP | ILM entry installed by MPLS-TP. |

The following table explains the show command output fields.

Table 62. show mpls ilm-table output field

| Field | Description |
|--------------|--|
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| ILM-ID | ILM identifier for the session. |
| LSP-Type | LSP type associated with each interface being protected. |
| Out-Label | Label received from downstream neighbor for route. |
| Out-Intf | Short name of the physical interface through which traffic goes to the protected link. |
| In label | Displays the ingress (incoming interface) label for this segment. |
| In intf | Installed as a result of configuring an interface. |
| Nexthop | Displays the IP address of the next hop. |
| Num ILMs | Number of ILMs in the session. |
| Swap Entries | Number of packets in the entry. |
| Pop Entries | Number of POP entries. |
| Upstream | Number of Upstream peer the FEC has been advertised. |

show mpls in-segment-table

Use this command to display detailed information about all entries in the Incoming Label Map (also known as in-segment) table.

Command Syntax

```
show mpls in-segment-table
```

Parameters

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls in-segment-table
  Owner: RSVP,#of pops: 1, fec: 192.168.0.5/32
    RX bytes:0, pkts:0, TX bytes:0, Swapped pkts:0, Popped pkts:0
LSP Type: ELSP_CONFIG
Class_Exp mapping:
CLASS      DSCP_value      EXP_value
  be        000000          0
  Cross connect ix: 1, in intf: eth0 in label: 52480 out-segment ix: 1
    Owner: RSVP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 1, owner: RSVP, out intf: eth1, out label: 52480
    Nexthop addr: 20.30.0.3      cross connect ix: 1, op code: Swap
  Cross connect ix: 1, in intf: eth0 in label: 52480 out-segment ix: 2
    Owner: RSVP, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 2, owner: RSVP, out intf: eth2, out label: 52481
    Nexthop addr: 30.30.0.3      cross connect ix: 1, op code: Swap
#
```

The following table explains the show command output fields.

Table 63. show mpls in-segment-table output field

| Field | Description |
|--------------|--|
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| RX bytes | Size of hello packets received from the neighbor. |
| Pkts | Number packet in the interface. |
| TX bytes | Size of the packets that transmitted to the neighbor. |
| Swapped pkts | Number of swapped packets in session. |

Table 63. show mpls in-segment-table output field (continued)

| Field | Description |
|------------------|--|
| Popped pkts | Number of popped packets in the interface. |
| LSP-Type | LSP type associated with each interface being protected. |
| CLASS | Creates a class map to be used for matching traffic to a specified class, and enters class-map configuration mode. |
| DSCP value | The value of the DSCP and DSCP classifier is used for routing Layer 3 packets. |
| EXP value | Sets the value of the MPLS EXP field on all imposed label entries. |
| Cross-connect ix | Displays the table index for the cross-connect. |
| Out-Label | Label received from downstream neighbor for route. |
| Out-Intf | Short name of the physical interface through which traffic goes to the protected link. |
| In label | Displays the ingress (incoming interface) label for this segment. |
| In intf | Installed as a result of configuring an interface. |
| Nexthop | Displays the IP address of the next hop. |
| Out-segment ix | Displays the outbound segment index. |
| Persistent | Displays whether the tunnel is persistent – Yes or No. |
| Admin Status | Indicates whether the user can administratively disable a peer while still preserving its configuration. Up = Yes, Down = No. |
| Oper Status | Displays the current status of the cross-connect segment – Up or Down. |
| Op code | PUSH = Replace the top label with another and then push one or more additional labels onto the label stack. SET = Set the next hop label. |

show mpls l2-circuit

Use this command to view MPLS-TP L2 circuit parameters.

Command Syntax

```
show mpls l2-circuit (detail|)
show mpls l2-circuit NAME (detail|)
```

Parameters

detail

Show detailed information

NAME

The name of the virtual circuit

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls l2-circuit detail
MPLS Layer-2 Virtual Circuit: vc1, id: 1 PW-INDEX: 1 service-tpid: 8100

Endpoint: 1.1.1.1
Control Word: 0
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe41
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active
Service-template : C1
Match criteria : 10-14, 16-20
```

The following table explains the show command output fields.

Table 64. show mpls l2-circuit output field

| Field | Description |
|------------------------------|--|
| MPLS Layer-2 Virtual Circuit | The MPLS virtual circuit on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches. |
| Endpoint | Endpoint address. |
| Control Word | Number of control words. |
| MPLS Layer-2 Virtual | The MPLS virtual circuit group on the egress PE router or switch and the specified |

Table 64. show mpls l2-circuit output field (continued)

| Field | Description |
|----------------------|---|
| Circuit Group | neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches. |
| Bound to interface | A bound service is the server in a client-server interface. |
| Virtual Circuit Type | Type of virtual circuit in the interface. |
| Service-template | Service Templates provides a powerful mechanism to configure advanced service-related options. |
| Match criteria | The match criteria under which redistribution is allowed for the current route-map. |

show mpls l2-circuit statistics

Use this command to display MPLS traffic statistics for L2 circuit.

Command Syntax

```
show mpls l2-circuit NAME statistics (access-port|network-port|)
```

Parameters

NAME

Name of L2 circuit

access-port

Displays the access port statistics

network-port

Displays the network port statistics

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls l2-circuit t1 statistics
MPLS Layer-2 Virtual Circuit: t1, id 100      # Virtual circuit name and ID
Access port statistics:
  RX: Input packets : 1000
     Input bytes   : 120000
  TX: Output packets : 0
     Output bytes  : 0
Network port statistics:
  RX: Input packets : 0
     Input bytes   : 0
  TX: Output packets : 1000
     Output bytes  : 120000
```

The following explains the show command output fields.

Table 65. show mpls l2-circuit statistics output field

| Field | Description |
|------------------------------|--|
| MPLS Layer-2 Virtual Circuit | The MPLS virtual circuit on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE routers or switches. |
| Virtual circuit name and ID | The MPLS virtual circuit identifier on the egress PE router or switch and the specified neighbor, testing the integrity of the Layer 2 circuit between the ingress and egress PE |

Table 65. show mpls l2-circuit statistics output field (continued)

| Field | Description |
|-------------------------|---|
| | routers or switches. |
| Access port statistics | Traffic statistics on Access port of VC/VPLS. |
| Network port statistics | Traffic statistics on Provider port of VC/VPLS. |
| RX | Number of received packets. |
| Input packets | Number of hello packets received from the neighbor. |
| Input bytes | Size of hello packets received from the neighbor. |
| TX | Number of packets transmitted. |
| Output packets | Number of hello packets sent to the neighbor. |
| Output bytes | Size of hello packets sent to the neighbor. |

show mpls mapped-routes

Use this command to view MPLS mapped routes.

Use the `no` parameter with this command to reset this configuration.

Command Syntax

```
show mpls mapped-routes
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls mapped-routes

Mapped-route      IPv4 FEC          MPLS-TP Tunnel
14.1.2.3/32       N/A              NH4
```

The following table explains the show command output fields.

Table 66. show mpls mapped-routes output field

| Field | Description |
|----------------|--|
| Mapped-route | Map the route of the interface. |
| IPv4 | IPv4 address of the neighbor interface. |
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| MPLS-TP Tunnel | MPLS-TP tunnel can be provisioned between two arbitrary nodes in an MPLS-TP enabled network. |

show mpls out-segment-table

Use this command to display detailed information of out-segment entries (also known as NHLFE) table.

Command Syntax

```
show mpls out-segment-table
```

Parameters

None

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls out-segment-table
  Out-segment with ix: 1, owner: RSVP, out intf: eth1, out label: 52480
    Nexthop addr: 20.30.0.3      cross connect ix: 1, op code: Swap
    TX bytes:0, pkts:0, error pkts:0, discard pkts:0

  Out-segment with ix: 2, owner: RSVP, out intf: eth2, out label: 52481
    Nexthop addr: 30.30.0.3      cross connect ix: 1, op code: Swap
    TX bytes:0, pkts:0, error pkts:0, discard pkts:0Zx
```

The following table explains the show command output fields.

Table 67. show mpls out-segment-table output field

| Field | Description |
|------------------|--|
| Out-segment ix | Displays the outbound segment index. |
| Out-Label | Label received from downstream neighbor for route. |
| Out-Intf | Short name of the physical interface through which traffic goes to the protected link. |
| Nexthop addr | Displays the IP address of the next hop. |
| Cross-connect ix | Displays the table index for the cross-connect. |
| Op code | PUSH = Replace the top label with another and then push one or more additional labels onto the label stack. SET = Set the next hop label. |
| Pkts | Number packet in the interface. |
| TX bytes | Size of the packets that transmitted to the neighbor. |
| Error pkts | Number of error packets. |

Table 67. show mpls out-segment-table output field (continued)

| Field | Description |
|--------------|---|
| Discard pkts | Number of packets discarded in the interface. |

show mpls owner-rib-table

This command is modified to display backup information when LDP_LFA is enabled.

Command Syntax

```
show mpls owner-rib-table
```

Parameters

None

Example

If LFA flag is disabled.

```
#show mpls owner-rib-table count
=====
RIB INFO  Type : LDP
Session IP Address: 41.41.41.41  Label Space: 0
Ftn Count 4, ILM Count 2 Dep ILM 2 VC 0
=====
```

If LFA is enabled and Backup entries is not present.

```
#show mpls owner-rib-table count
=====
RIB INFO  Type : LDP
Session IP Address: 41.41.41.41  Label Space: 0
Ftn Count 4, ILM Count 2 Dep ILM 2 VC 0
Backup Info
Ftn Count 0, ILM Count 0 Dep ILM 0
=====
```

If Backup entries are present.

```
#show mpls owner-rib-table count
=====
RIB INFO  Type : LDP
Session IP Address: 41.41.41.41  Label Space: 0
Ftn Count 4, ILM Count 2 Dep ILM 2 VC 0
Backup Info
Ftn Count 4, ILM Count 0 Dep ILM 0
=====

RIB INFO  Type : LDP
Session IP Address: 17.17.17.17  Label Space: 0
Ftn Count 4, ILM Count 2 Dep ILM 2 VC 0
Backup Info
Ftn Count 4, ILM Count 0 Dep ILM 0
```

show mpls qos-resource

Use this command to display detailed QoS resource information.

Command Syntax

```
show mpls qos-resource IFNAME
```

Parameters

IFNAME

Display the interface name for a QoS resource

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls qos-resource eth1
<*****>
QOS RESERVED TABLE

<*****>
HOLD PRIORITY : 0

HOLD PRIORITY : 1

HOLD PRIORITY : 2

HOLD PRIORITY : 3

HOLD PRIORITY : 4

HOLD PRIORITY : 5

HOLD PRIORITY : 6

HOLD PRIORITY : 7
<*****>
  QOS AWAITING TABLE (static resources)
<*****>
HOLD PRIORITY : 0

HOLD PRIORITY : 1

HOLD PRIORITY : 2

HOLD PRIORITY : 3

HOLD PRIORITY : 4

HOLD PRIORITY : 5
```

```
HOLD PRIORITY : 6  
  
HOLD PRIORITY : 7  
TSUP-173>
```

The following table explains the show command output fields.

Table 68. show mpls qos-resource output fields

| Field | Description |
|---------------------------------------|---|
| QOS RESERVED TABLE | FTM/ILM entries for which QOS is reserved. |
| HOLD PRIORITY | Determines the degree to which an LSP holds onto its session reservation after the LSP has been set up successfully |
| QOS AWAITING TABLE (static resources) | FTN/ILM entries for which QOS reservation is pending. |

show mpls vc-table

Use this command view configured virtual circuit (VC) components

Command Syntax

```
show mpls vc-table
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show mpls vc-table
```

```
VC-ID  Vlan-ID  Inner-Vlan-ID  Access-Intf  Network-Intf  Out Label  Tunnel-Label  Nexthop  Status
500    N/A       N/A           eth2         eth1          544        57           N/A      Active
#
```

The following table explains the show command output fields.

Table 69. show mpls vc-table output fields

| Field | Description |
|----------------|--|
| VC-ID | The virtual circuit ID for the Provider Edge (PE) MPLS. |
| Vlan-ID | Virtual LAN (VLAN) ID number. |
| Inner-Vlan-ID | Inner Virtual LAN (VLAN) ID number. |
| Access-Intf | The Interface Access page provides a method with which to control access to specific areas of the interface. |
| Network-Intf | A networking interface allows a computer or mobile device to connect to a local area network (LAN) using Ethernet as the transmission mechanism. |
| Out Label | Label received from downstream neighbor for route. |
| Tunnel-Label | Used to provide reachability between PE devices. |
| Nexthop Status | Displays the network status of the next hop. |

show mpls vrf

Use this command to display detailed information of all the configured VRF entries. Specify the name of the VRF to display information about a specific VRF entry.

Command Syntax

```
show mpls vrf-table
show mpls vrf-table VRFNAME (count|)
```

Parameters

VRFNAME

Display the MPLS VRF table by its configured name

count

Display the MPLS VRF FTN's count

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 1.3. Enhanced the command for ECMP in OcNOS version 6.6.1

Examples

```
1. 15 routes having primary and backup paths
2. 10 routes having primary paths (no backup)
#show mpls vrf new_vrf count
-----
Num VRF-FTNs           : 25           [UP: 25, INSTALLED: 25]
-----
Num IPV6 VRF-FTNs      : 0           [UP: 0, INSTALLED: 0]
-----
```

The following explains the show command output fields.

Table 70. show vrf-table output fields

| Field | Description |
|-------------------|--|
| Num VRF-FTNs | Number of FEC-to-NHLFE map counters in VRF protocol. |
| Num IPV6 VRF-FTNs | Number of VRF FEC-to-NHLFE map counters in protocol. |

show mpls vrf-forwarding-table vrf

This CLI can be used to display a tabular output of the VRF forwarding entries received from the remote PE via MPBGP.

Command Syntax

```
show mpls vrf-forwarding-table vrf <VRFNAME>
```

Parameters

VRFNAME

Display the MPLS VRF table by its configured name

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS version 4.1. Enhanced the command for ECMP in OcNOS version 6.6.1

Examples

```
#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services
```

| Code | FEC | FTN-ID | VRF-ID | Nhlfe-ID | Pri | Out-Label | Out- |
|------|-----------------|--------|-----------|----------|-----|-----------|------|
| Intf | Nexthop | UpTime | Ext-Color | | | | |
| B> | 198.1.1.0/24 | 1 | 3 | 41 | (e) | - | - |
| | 00:04:40 | - | | 38 | Yes | 21 | (D) |
| | 2.2.2.2 | - | | 36 | Yes | 20 | (D) |
| | 2.2.2.2 | - | | 13 | Yes | 30 | - |
| | 3.3.3.3 | - | | 14 | Yes | 40 | - |
| | 4.4.4.4 | - | | 40 | Yes | 50 | - |
| | 5.5.5.5 | - | | 55 | Yes | 61 | - |
| | 222.222.222.222 | - | | 31 | Yes | 61 | - |
| B> | 2222::/64 | 2 | 3 | 47 | (e) | - | - |
| | 00:04:40 | - | | 44 | Yes | 23 | (D) |
| | 2.2.2.2 | - | | 42 | Yes | 22 | (D) |
| | 2.2.2.2 | - | | 19 | Yes | 31 | - |
| | 3.3.3.3 | - | | 17 | Yes | 41 | - |
| | 4.4.4.4 | - | | 46 | Yes | 51 | - |

```

5.5.5.5      -      55
              33      Yes  63      -
222.222.222.222 -      -
UFI-28SMT#
UFI-28SMT#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code   FEC           FTN-ID VRF-ID   Nhlfe-ID   Pri   Out-Label   Out-
Intf    Nexthop          UpTime    Ext-Color
B> 198.1.1.0/24    1        3        41        (e)   -           -
    00:05:03      -
    2.2.2.2        -        -        38        Yes   21          -
    2.2.2.2        -        -        36        Yes   20          -
    2.2.2.2        -        -        13        Yes   30          -
    3.3.3.3        -        -        14        Yes   40          -
    4.4.4.4        -        -        40        Yes   50          -
    5.5.5.5        -        55        31        Yes   61          -
    222.222.222.222 -        -        47        (e)   -           -
B> 2222::/64      2        3        44        Yes   23          -
    00:05:03      -
    2.2.2.2        -        -        42        Yes   22          -
    2.2.2.2        -        -        19        Yes   31          -
    3.3.3.3        -        -        17        Yes   41          -
    4.4.4.4        -        -        46        Yes   51          -
    5.5.5.5        -        55        33        Yes   63          -
    222.222.222.222 -        -
UFI-28SMT#
UFI-28SMT#
0          28          xe4          10.143.73.6

```

The following table explains the show command output fields.

Table 71. show mpls vrf-forwarding-table vrf output fields

| Field | Description |
|-------------|--|
| Owner | Displays the creator of this entry, typically a protocol such as BGP. |
| FEC | Displays the Forward Equivalency Class (FEC) for this entry. |
| FTN-ID | FEC-to-NHLFE identification. |
| Oper-Status | Displays the current status of the entry – Up or Down. It will be “UP” if the vrf entry is installed in the forwarder and it will be in “DOWN” state if the vrf entry is not installed in the forwarder. |
| Out-Label | Displays the egress label for this FTN. |
| Tunnel-id | Tunnel identification to which packets of this FTN are going. |

| Field | Description |
|-----------|--|
| NHLFE-id | Next Hop Label Forwarding Entry identification (also known as out-segment entry identification). |
| Out-Intf | Name of the physical interface through which traffic goes. |
| Nexthop | Displays the IP address of the next hop. |
| Ext-Color | Extended-community color advertised by BGP |

show running-config interface mpls

Use this command to show the running system status and configuration for an MPLS interface.

Command Syntax

```
show running-config interface IFNAME mpls
```

Parameters

IFNAME

Display information for this interface name

Command Mode

Privileged execution mode and Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show running-config interface eth1 mpls
```

show running-config mpls

Use this command to show any Multi-Protocol Label Switching (MPLS) related running configuration.

Command Syntax

```
show running-config mpls
```

Parameters

None

Command Mode

Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
>enable
#show running-config mpls
!
mpls propagate-ttl
!
!
!
```

show running-config vc

Use this command to show any Virtual Private Wire Service (VPWS) related running configuration.

Command Syntax

```
show running-config vc
```

Parameters

None

Command Mode

Privileged execution mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
#show running-config vc
!
mpls l2-circuit vc1 1 2.2.2.2
!
mpls l2-circuit vc2 3 2.2.2.2
  tunnel-select-policy p1
!
!
interface xe2
  mpls-l2-circuit vc1 service-template s1 primary
  mpls-l2-circuit vc2 service-template s3 primary
!
```

show running-config vpls

Use this command to show any Virtual Private LAN Service (VPLS) related running configuration.

Command Syntax

```
show running-config vpls
```

Parameters

None

Command Mode

Privileged execution mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
#show running-config vpls
!
mpls vpls vpls1 2
  signaling ldp
  vpls-type vlan
  vpls-peer 2.2.2.2
  exit-signaling
  exit-vpls
!
mpls vpls vpls2 4
  signaling ldp
  vpls-type vlan
  vpls-peer 2.2.2.2 tunnel-select-policy p1
  exit-signaling
  exit-vpls
!
!
interface xe2
  mpls-vpls vpls1 service-template s2
  mpls-vpls vpls2 service-template s4 split-horizon group access1
```

show vccv statistics

Use this command to display VCCV messages received prior to advertising capability.

Command Syntax

```
show vccv statistics
```

Parameters

None

Command Mode

Privilege mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following is the sample output for `show vccv statistics` command.

```
#show vccv statistics
CC Mismatch Discards - 10
```

The following table explains the show command output fields.

Table 72. show vccv statistics output fields

| Field | Description |
|----------------------|---|
| CC Mismatch Discards | Number of CC mismatch packets received from neighbor discarded. |

srlg-disjoint

Use this command to set how to avoid the Shared Risk Link Groups (SRLGs) of a protected primary.

A fast-reroute or secondary path for an LSP that is disjoint from the primary ensures that a single point of failure on a particular link does not bring down both the primary and secondary paths in the LSP.



Notes:

- The SRLG option configured in RSVP-TRUNK mode (see the [secondary-priority srlg-disjoint \(page 1434\)](#) command) takes higher preference than the option configured in RSVP router mode (this command).
- For one-to-one protection, when both “[primary fast-reroute node-protection](#)” and “`srlg-disjoint preferred`” commands are configured, detour path selection occurs in the following order:
 - Find a path for “Node protection along with SRLG protection”; if no path is found, then
 - Find a path for “Link protection along with SRLG protection”; if no path is found, then
 - Find a path for “Node protection without SRLG protection”; if no path is found, then
 - Find a path for “Link protection without SRLG protection”.
- When SRLG disjointness is configured on a PLR, auto-bypass tunnels are created using the SRLG constraints of the outgoing interface for the protected LSP. If the SRLG values on that interface are updated, a new auto-bypass tunnel reflecting the revised SRLG set is created and used to protect the LSPs. The previous auto-bypass tunnel is marked invalid and is removed after all protected LSP references to that tunnel are cleared.

Use the `no` form of this command to not avoid the SRLGs of a protected interface.

Command Syntax

```
srlg-disjoint (forced|preferred)
no srlg-disjoint
```

Parameters

forced

The router does not create the secondary or backup tunnel unless it avoids SRLGs of the primary-path/protected-interface.

preferred

With two explicit paths, the first explicit path tries to avoid the SRLGs of the primary-path/ protected interface. If that does not work, the secondary or backup tunnel uses the second path (which ignores SRLGs).

Default

None

Command Mode

Router RSVP mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#router rsvp
(config-rsvp)# srlg-disjoint forced
```


trace mpls

Use this command to isolate the failure point of an LSP. It is used for hop-by-hop fault localization and path tracing. The MPLS LSP traceroute feature depends on the expiration of the packet's Time to Live (TTL) value that carries the echo request.

Users can initiate a Segment Routing traceroute operation only when the Segment Routing control plane is available at the origin. Use this command to verify connectivity when an LSP traverses entirely through a Segment Routing domain.

Command Syntax

```
trace mpls (ospf-sr | isis-sr) (ipv4 A.B.C.D/M | sr-policy SR_POLICY_NAME candidate-path CANDIDATE_
PATH_ID) ({reply-mode (2)|flags | destination A.B.C.D
| source A.B.C.D | ttl <1-255> | timeout <1-500> | detail|})
```

Parameters

ospf-sr

Specifies the OSPF Segment Routing protocol domain.

isis-sr

Specifies the IS-IS Segment Routing protocol domain.

ipv4 A.B.C.D/M

Specifies the IPv4 prefix (FEC) of the LSP to trace.

sr-policy SR_POLICY_NAME

Specifies the Segment Routing Policy name for which the LSP traceroute is initiated.

candidate-path CANDIDATE_PATH_ID

Specifies the dynamic candidate path in the SR policy with the given candidate path ID.

reply-mode (2)

Defines the MPLS echo reply format. 2 - Reply via UDP or IP packet (default mode).

flags

Validate FEC stack.

destination A.B.C.D

Sets the destination IPv4 address used in the echo request packet.

source A.B.C.D

Sets the source IPv4 address used in the echo request packet.

ttl <1-255>

Specifies the Time To Live (TTL) value for the probe packet. Controls how many hops the trace explores. The range is 1–255. The default TTL is 255.

timeout <1-500>

Specifies the time to wait before declaring a probe as failed, in seconds. The range is 1–500. The default timeout is 60 seconds.

detail

Displays detailed output for the trace probe.

Default

None

Command Mode

Privileged Exec mode

Applicability

- Introduced in OcNOS Version SP 4.0.
- Removed repeat, interval, and force-explicit-null parameters in OcNOS version 6.6.1.

Examples

The following example shows how to trace an MPLS Label Switched Path (LSP) within an OSPF Segment Routing domain to verify hop-by-hop connectivity and identify the point of failure.

```
#trace mpls ospf-sr ipv4 10.0.1.33/32 destination 127.1.0.1 source 10.0.1.11 detail
Tracing MPLS Label Switched Path to 2.32.0.0, timeout is 5 seconds
```

Codes:

```
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed
```

Type 'Ctrl+C' to abort

```
0 10.11.1.1 [Labels: 16500]
R 1 10.0.1.1 [Labels: 16500] 0.97 ms
! 2 10.0.1.33 3.90 ms
```

```
# trace mpls isis-sr ipv4 2.2.2.2/32 detail
# trace mpls ospf-sr ipv4 2.2.2.2/32 detail
```

```
# trace mpls isis-sr sr-policy aaa candidate-path 1
# trace mpls ospf-sr sr-policy ABC candidate-path 1
```

suppress-oper-log mpls l2vpn

Use this command to suppress the operator logs for the MPLS L2VPN service.

Use the `no` parameter with this command to display the operator logs for the MPLS L2VPN service.

Command Syntax

```
suppress-oper-log mpls l2vpn  
no suppress-oper-log mpls l2vpn
```

Parameters

l2vpn

Operator logs of L2VPN.

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS#configure terminal  
OcNOS (config)#suppress-oper-log mpls l2vpn  
OcNOS (config)#commit
```

trace mpls

Use this command to trace the route traversed by a specified echo request packet in an MPLS protocol. Trace requests can be configured for LDP, RSVP, L2 VC, VPLS, Segment Routing and L3 VPN label switched paths.



Note: When the tunnel uses the bypass LSP, the "trace mpls rsvp" command does not show bypass path hop details (shows timeout).

Command Syntax

```
trace mpls (6pe default X:X::X:X/M|ldp A.B.C.D/M|rsvp (tunnel-name NAME|egress A.B.C.D)|l3vpn VRFNAME
A.B.C.D/M|ipv4 A.B.C.D/M) ({reply-mode (2)|flags|destination A.B.C.D|source A.B.C.D|timeout <1-
500>|force-explicit-null|detail})
```

Parameters

6pe

FEC type is 6pe

default

VPN Instance Name (default)

X:X::X:X/M

6pe prefix address

ldp

FEC type is LDP

A.B.C.D/M

LDP prefix address

rsvp

FEC type is RSVP

tunnel-name

RSVP tunnel name

NAME

Tunnel name string

egress

RSVP tunnel egress

A.B.C.D

RSVP tunnel egress address

l3vpn

FEC type is MPLS VPN (L3-VPN)

VRFNAME

VPN instance name

A.B.C.D./M

VPN prefix

ipv4

FEC type generic; use for static/SNMP label switched paths

A.B.C.D/M

IPv4 prefix address

X:X::X:X/M

VPNv6 prefix

Default

60 seconds.

Command Mode

Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#trace mpls ipv4 10.10.0.0/24 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 timeout 65
detail force-explicit-null

#trace mpls l3vpn vrfa 10.10.0.0/24 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 timeout
65 detail force-explicit-null

#trace mpls ldp 10.10.0.0/24 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 timeout 65
detail force-explicit-null

#trace mpls rsvp egress 1.2.3.5 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 timeout 65
detail force-explicit-null

#trace mpls rsvp tunnel-name tun1 reply-mode 2 flags destination 127.1.2.3 source 10.10.0.1 timeout
65 detail force-explicit-null
```

tunnel-name

Use this command to configure tunnel name for the MPLS transport tunnel to be used for the MPLS layer-2 virtual circuit.

Use the no parameter with this command to delete tunnel name from the MPLS layer-2 virtual circuit.

Command Syntax

```
tunnel-name NAME  
no tunnel-name
```

Parameters

NAME

Identifying name for MPLS Tunnel

Default

Disabled

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 1.3

Example

```
#configure terminal  
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4  
(config-pseudowire)#tunnel-name pe1-to-pe2
```

tunnel-id

Use this command to configure tunnel identifier for the MPLS transport tunnel to be used for the MPLS layer-2 virtual circuit.

Use the no parameter with this command to delete tunnel identifier from the MPLS layer-2 virtual circuit.

Command Syntax

```
tunnel-id <1-5000>  
no tunnel-id
```

Parameters

<1-5000>

Identifying value for Tunnel-id

Default

Disabled

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 1.3

Example

```
#configure terminal  
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4  
(config-pseudowire)#tunnel-id 22
```

tunnel-select-policy

Use this command to configure tunnel selection policy name for the MPLS transport tunnel to be used for the MPLS layer-2 virtual circuit.

Use the no parameter with this command to delete tunnel selection policy name from the MPLS layer-2 virtual circuit.

Command Syntax

```
tunnel-select-policy POLICYNAME  
no tunnel-select-policy
```

Parameters

POLICYNAME

Selection policy name for MPLS Tunnel

Default

Disabled

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 1.3

Example

```
#configure terminal  
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4  
(config-pseudowire)#tunnel-select-policy policy1
```

vccv cc-type

Use this command to configure the VCCV control channel for MPLS layer-2 virtual circuit.

Use the no parameter with this command to disable control channel from MPLS layer-2 virtual circuit.

Command Syntax

```
vccv cc-type (type-1|type-2|type-3)
no vccv cc-type (type-1|type-2|type-3)
```

Parameters

type-1

CC Type 1 - PWE3 Control Word with 0001b as first nibble

type-2

CC Type 2 - MPLS Router Alert Label

type-3

CC Type 3 - MPLS PW Label with TTL == 1

Default

Disabled

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4
(config-pseudowire)# vccv cc-type type-2
```

vccv cv-type

Use this command to configure the VCCV control verification for MPLS layer-2 virtual circuit.

Use the no parameter with this command to disable control verification from MPLS layer-2 virtual circuit.

Command Syntax

```
vccv cv-type (type-1|type-2|type-3|type-4)
no vccv cv-type (type-1|type-2|type-3|type-4)
```

Parameters

type-1

BFD IP/UDP-encapsulated for PW Fault Detection only

type-2

BFD IP/UDP-encapsulated for PW Fault Detection and AC/PW Fault Status Signalling

type-3

BFD PW-ACH-encapsulated for PW Fault Detection only

type-4

BFD PW-ACH-encapsulated for PW Fault Detection and AC/PW Fault Status Signalling

Default

Disabled

Command Mode

Configure Pseudowire mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#mpls l2-circuit mycircuit 45678 1.2.3.4
(config-pseudowire)# vccv cv-type type-1
```

Virtual Private LAN Service Commands

| | |
|--|------|
| ac-admin-status | 1496 |
| ac-description | 1497 |
| clear mpls vpls | 1498 |
| control-word | 1499 |
| exit-signaling | 1500 |
| exit-if-vpls | 1501 |
| learning disable (VPLS Mode) | 1502 |
| learning disable (Interface VPLS Mode) | 1503 |
| learning enable | 1504 |
| mpls vpls | 1505 |
| mpls-vpls service-template | 1506 |
| no learning | 1507 |
| show mpls vpls | 1508 |
| show mpls vpls mac-address | 1516 |
| show mpls vpls split-horizon | 1518 |
| show mpls vpls statistics | 1520 |
| signaling ldp | 1522 |
| signaling bgp | 1523 |
| static-mac | 1524 |
| tunnel-select-policy | 1525 |
| ve-id | 1526 |
| vpls-ac-group | 1527 |
| vpls-description | 1528 |
| vpls fib-entry | 1529 |
| vpls-mtu | 1531 |
| vpls-peer | 1532 |
| vpls-peer manual | 1534 |
| vpls-type | 1535 |
| vpls-vc | 1536 |
| vpls-route-map | 1537 |

ac-admin-status

Use this command to configure the admin status of an attachment circuit specific to a VPLS instance.

Command Syntax

```
ac-admin-status down  
no ac-admin-status
```

Parameter

down

set the admin role as DOWN

Default

By default, ac admin status is up.

Command Mode

Interface VPLS

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows the configuration of admin status for attachment circuit specific to VPLS instance

```
#configure terminal  
(config)#interface xel  
(config-if)#mpls-vpls vpls1 service-template st1  
(config-if-vpls)#no ac-admin-status
```

ac-description

Use this command to add description for an attachment circuit specific for a VPLS instance

Use the no parameter with this command to remove the description

Command Syntax

```
ac-description LINE
```

Parameter

LINE

Characters describing this AC

Default

By default, ac description LINE is disabled.

Command Mode

Interface VPLS

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows the configuration of description for attachment circuit specific to VPLS instance

```
#configure terminal
(config)#interface xel
(config-if)#mpls-vpls vpls1 service-template st1
(config-if-vpls)#ac-description AC1_VPLS1
```

clear mpls vpls

Use this command to clear VPLS data.

Command Syntax

```
clear mpls vpls (NAME |) mac-addresses
```

Parameters

NAME

Clear data for the VPLS instance with name given

mac-addresses

Flush all MAC addresses for a VPLS instance

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#clear mpls vpls VPLS_123 mac-addresses
```

control-word

Use this command to enable the control-word for a VPLS instance.

Use the `no` parameter with this command to disable the control-word.



Notes:

- Control-word (CW) negotiation can be `Preferred` or `Not Preferred`. OcNOS follows the `Not Preferred` mode. OcNOS does not rely on whether the peer has control-word enabled.
 - If both sides enable control-word, OcNOS installs the negotiated CW as 1.
 - If only one side enables control-word, OcNOS installs the negotiated CW as 0.
 - In either case, OcNOS does not bring the Virtual Circuit (VC) down.
- Some interoperability devices use the `Preferred` mode in control-word negotiation. When these devices are configured with control-word, they expect the other end device (OcNOS) to also have control-word configured. If the OcNOS does not match its configuration, the interoperability devices set the PW status to `Not Forwarding` or raise a fault. As a result, the VC on OcNOS appears as `Down` because the remote pseudowire (PW) status indicates a fault.

Command Syntax

```
control-word
no control-word
```

Parameters

None

Default

Disabled

Command Mode

VPLS configuration mode

Applicability

Introduced in OcNOS version 4.1.

Example

```
(config)#mpls vpls ldp-vpls11 11
(config-vpls)#control-word

(config-vpls)#no control-word
(config-vpls)#exit-vpls
```

exit-signaling

Use this command to exit the VPLS signaling configuration mode, and start signaling. To configure signaling with LDP, see the [signaling ldp \(page 1522\)](#) command. Other VPLS signaling configuration commands include [show mpls vpls \(page 1508\)](#), [show mpls vpls \(page 1508\)](#), [vpls-vc \(page 1536\)](#), [vpls-ac-group \(page 1527\)](#), and [vpls-peer \(page 1532\)](#).



Note: It is **critical** to give this command after all VPLS signaling configurations are complete, otherwise signaling does not start.

Command Syntax

```
exit-signaling
```

Parameters

None

Default

No default value is specified.

Command Mode

VPLS signal mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
# configure terminal
(config)#mpls vpls test 100
(config-vpls)#signaling ldp
(config-vpls-sig)#exit-signaling
```


exit-if-vpls

Use this command to exit from Interface VPLS mode

Command Syntax

```
exit-if-vpls
```

Parameter

None

Default

No default value is specified.

Command Mode

Interface VPLS

Applicability

This command was introduced before OcNOS version 1.3.

Example

This example shows exiting from interface VPLS mode

```
#configure terminal
(config)#interface xel
(config-if)#mpls-vpls vpls1 service-template st1
(config-if-vpls)#ac-description AC1_VPLS1
(config-if-vpls)#exit-if-vpls
(config-if-vpls)#exit
```

learning disable (VPLS Mode)

Use this command to disable learning for a VPLS instance.

Use the `no` form of this command to enable learning on a VPLS instance.



Note: This command disables learning on all the attachment circuits and pseudo-wires belonging to that VPLS instance.

Command Syntax

```
learning disable
no learning disable
```

Parameter

None

Default

By default, learning disable is disabled.

Command Mode

VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config-vpls)#mpls vpls vpls2 vlan 3
(config-vpls)#learning disable
(config-vpls)#exit

#configure terminal
(config-vpls)#mpls vpls vpls2 vlan 3
(config-vpls)#no learning disable
(config-vpls)#exit
```

learning disable (Interface VPLS Mode)

Use this command to disable learning on a particular Attachment Circuit (AC) interface.

Use the [learning enable \(page 1504\)](#) command to enable learning on a particular AC interface.



Note: This command disables MAC learning only on that interface.

Command Syntax

```
learning disable
```

Parameter

None

Default

By default, learning disable is disabled.

Command Mode

Interface VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#mpls-vpls vpls1 service-template st1
(config-if-vpls)#learning disable
(config-if-vpls)#exit
```

learning enable

Use this command to enable learning on a particular attachment circuit (AC) interface.

Use the [learning disable \(Interface VPLS Mode\) \(page 1503\)](#) command to disable learning on a particular AC interface.



Note: This command enables MAC learning only on that AC interface.

Command Syntax

```
learning enable
```

Parameter

None

Default

By default, learning enable is enabled.

Command Mode

Interface VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#mpls-vpls vpls1 service-template st1
(config-if-vpls)#learning enable
(config-if-vpls)#exit
```

mpls vpls

Use this command to create an instance of MPLS-based Virtual Private LAN Services (VPLS).

Use the `no` parameter with this command to delete an MPLS-based VPLS instance.



Note: OcNOS supports configuring the same VPLS instance identifier (VPLS-ID) as the virtual circuit identifier (VC-ID) with different LDP neighbors.

Command Syntax

```
mpls vpls NAME
mpls vpls NAME <1-4294967295>
no mpls vpls NAME
```

Parameters

NAME

VPLS instance identifier

<1-4294967295>

VPLS instance identifier

Default

By default, mpls vpls is disabled.

Command Mode

Configuration mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#mpls vpls t1 6489
(config-vpls)#exit
```

mpls-vpls service-template

Use this command to bind a VPLS instance to a service template.

Use the no parameter with this command unbind the VPLS instance from service template.

Command Syntax

```
mpls-vpls VPLS_NAME service-template TEMPLATE_NAME
no mpls-vpls VPLS_NAME service-template TEMPLATE_NAME
```

Parameters

VPLS_NAME

VPLS instance name

TEMPLATE_NAME

Service template name

Default

None

Command Mode

Interface mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth1
(config-if)#switchport
(config-if)#mpls-vpls VPLS1 service-template C1
(config-if-vpls)#exit-if-vpls

(config-if)#no mpls-vpls VPLS1 service-template C1
(config-if)#exit-if-vpls
```

no learning

Use this command to reset learning on a particular AC-interface to the global learning configuration.

Command Syntax

```
no learning
```

Parameter

None

Default

By default, no learning is disabled.

Command Mode

Interface VPLS mode and VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal
(config)#interface eth0
(config-if)#mpls-vpls vpls1 service-template st1
(config-if-vpls)#no learning
(config-if-vpls)#exit
(config)#

#configure terminal
(config)#mpls vpls vpls5 vlan 34
(config-vpls)#learning limit 500
(config-vpls)#exit
(config)#
```

show mpls vpls

Use this command to display logging information configured for MPLS.

Command Syntax

```
show mpls vpls
show mpls vpls debug-detail
show mpls vpls detail
show mpls vpls mesh
show mpls vpls NAME
show mpls vpls NAME mesh
show mpls vpls NAME spoke
show mpls vpls spoke
show mpls vpls count
```

Parameters

debug-detail

Show detailed VPLS information for debugging purpose

detail

Display detailed VPLS information

mesh

Display MPLS VPLS Mesh Forwarding information. Use this parameter to display information about all core Virtual Circuit (VC) connections for all VPLS instances. Give the name of a VPLS instance to display information about that instance.

NAME

Display the identifying string for the VPLS domain

spoke

Display MPLS VPLS Spoke Forwarding information. Use this parameter to display information about all spoke VC connections for all VPLS instances. Give the name of a VPLS instance to display information about that instance.

count

Display the count of VPLS instances

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

Using `show mpls vpls` command without parameters displays information about all VPLS instances.

The example below displays information about the VPLS instance `v1`, returned when using the `NAME` parameter.

```
#show mpls vpls t1
Virtual Private LAN Service Instance: t1, ID: 1
Group ID: 0, VPLS Type: Ethernet VPLS, Configured MTU: 0
```



```

Description: none
Configured interfaces: none
Mesh Peers: 192.168.0.80 (Up)
            192.168.0.90 (Up)
Spoke Peers: t100 (Up)
#

```

The following table explains the show command output fields.

Table 73. show mpls vpls t1 output field

| Field | Description |
|--------------------------------------|---|
| Virtual Private LAN Service Instance | Number of VPLAN service instance. |
| ID | VPLAN identification detail for service instance. |
| Group ID | Group identification detail for VLAN. |
| VPLS Type | Type of VPLS in the interface. |
| Configured MTU | Number of configured MTU in the VPLs. |
| Description | Details of VPLS. |
| Configured interfaces | Description of the configured interfaces. |
| Mesh Peers | Configuring the VPLS mesh peers. |
| Spoke Peers | Configuring the VPLS spoke peers. |

The example below displays the name of the VPLS instance, its ID, they type of instance (Ethernet), the M and S peers, and the signaling protocol. For the first entry, the signaling protocol is BGP and for the second entry it is LDP.

```

#show mpls vpls
Name  VPLS-ID    Type           MPeers  SPeers  SIG-Protocol
v1    100        Ethernet       1        0       BGP
v3    300        Ethernet       1        0       LDP

```

The following table explains the show command output fields.

Table 74. show mpls vpls output field

| Field | Description |
|--------------|---|
| Name | Type of the MPLS protocol. |
| VPLS-ID | Identification detail of VPLS. |
| Type | Type of VPLS in MPLS protocol. |
| Mesh Peers | Configuring the VPLS mesh peers. |
| Spoke Peers | Configuring the VPLS spoke peers. |
| SIG-Protocol | Type of protocol in MPLS configuration. |

The example below displays the output when using the debug-detail parameter. It displays information for VPLS instance `vpls1`, including the signaling protocol, create time, uptime, total down time, associated transport, traffic statistics if non-zero and statistics profile enabled, and so on.

```

#show mpls vpls debug-detail
Virtual Private LAN Service Instance: vpls1, ID: 300
SIG-Protocol: LDP

```

```
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100

Configured interfaces:
Interface: xe4
Status: Up
Service-template : vpls1
Match criteria : 440

Mesh Peers:
  2.2.2.2 State: Installed
    CreateTime: 08:56:53, UpTime: 08:56:53
    TotalDownTime: N/A, FlapCount: 0, LastFlapReason: N/A
    Tunnel-Id: 11
    Transport: owner CLI, ftn_ix 9
  3.3.3.3 State: Installed
    CreateTime: 08:56:53, UpTime: 08:56:53
    TotalDownTime: N/A, FlapCount: 0, LastFlapReason: N/A
    Tunnel-Name: t1
    Transport: owner RSVP, ftn_ix 8
  4.4.4.4 State: Installed
    CreateTime: 08:56:53, UpTime: 08:49:39
    TotalDownTime: 00:01:11, FlapCount: 1, LastFlapReason: FIB delete from protocol
    Tunnel-Policy: p1
    Transport: owner SR_POLICY, ftn_ix 13
```

The example below displays the output when using the `detail` parameter. It displays information for VPLS instance `v1`, including the signaling protocol.

```
#show mpls vpls debug-detail
Virtual Private LAN Service Instance: vpls1, ID: 300
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100

Configured interfaces:
Interface: xe4
Status: Up
Service-template : vpls1
Match criteria : 440

Mesh Peers:
  2.2.2.2 State: Installed
    CreateTime: 08:56:53, UpTime: 08:56:53
    TotalDownTime: N/A, FlapCount: 0, LastFlapReason: N/A
    Tunnel-Id: 11
    Transport: owner CLI, ftn_ix 9
  3.3.3.3 State: Installed
    CreateTime: 08:56:53, UpTime: 08:56:53
    TotalDownTime: N/A, FlapCount: 0, LastFlapReason: N/A
    Tunnel-Name: t1
```

```

Transport: owner RSVP, ftn_ix 8
4.4.4.4 State: Installed
CreateTime: 08:56:53, UpTime: 08:49:39
TotalDownTime: 00:01:11, FlapCount: 1, LastFlapReason: FIB delete from protocol
Tunnel-Policy: pl
Transport: owner SR_POLICY, ftn_ix 13

```

The following table explains the show command output fields.

Table 75. show mpls vpls details output field

| Field | Description |
|--------------------------------------|---|
| Virtual Private LAN Service Instance | Number of VPLS service instance. |
| ID | VPLS identification detail for service instance. |
| SIG-Protocol | Type of protocol in MPLS configuration. |
| Attachment-Circuit | Details of the attached circuit in interface. |
| Learning | State of the interface. |
| Group ID | Group identification detail for VLAN. |
| VPLS Type | Type of VPLS in MPLS protocol. |
| Configured MTU | Number of configured MTU in the VPLs. |
| Description | Details of VPLS. |
| Service-tpid | Service TP identifier configured for the VPLS PW. |
| Operating mode | Type of mode in the interface. |
| Svlan Id | Configures a specific virtual LAN (VLAN). |
| Svlan Tpid | Service vlan TP identifier for the VPLS PW. |
| Redundancy admin role | Creating a Backup Administrator Role. |
| Redundancy oper role | Operational Role of the VPLS instance. |
| Configured interfaces | Details of the configured interfaces. |
| Interface | Selects an interface to configure. |
| Oper-state | Displays the current status of the cross-connect segment – Up or Down. |
| Service-template | Used to configure advanced service-related option. |
| Match criteria | Identifies prefix characteristics (network, BGP path attribute, nexthop, and so on) for a specific sequence. |
| Mesh Peers | Configuring the VPLS mesh peers. |
| PW Status Local | Used to perform limited local configuration changes, monitor device status and utilization, and simple local troubleshooting. |
| Remote | PW status of Remote end. |

The example below displays the output provided when using the `mesh` parameter without a specific VPLS name.

```

VPLS-ID  Peer Addr  Tunnel-Label  In-Label  Network-Intf  Out-Label  Lkps/St  PW-INDEX  SIG-Protocol
100      2.2.2.2    N/A          52503     eth2          53258     0/Dn     2         BGP
300      2.2.2.2    N/A          none      N/A          none      0/Dn     1         LDP

```

The following table explains the show command output fields.

Table 76. show mpls vpls output field

| Field | Description |
|--------------|---|
| VPLS-ID | Identification details of the VPLS. |
| Peer Addr | IP address of the peer device. |
| Tunnel-Label | Tunnel label used for the next segment. |
| In-label | Displays the ingress (incoming interface) label for this segment. |
| Out-Label | Label received from downstream neighbor for route. |
| Network-Intf | Installed as a result of configuring an interface. |
| Lkps/St | Opcode and Status of the VPLS PW. |
| PW-INDEX | Index of the VPLS entry in PW table. |
| SIG-Protocol | Signaling protocol used for VPLS labels advertisement. |

The following is a sample output of the `show mpls vpls detail` command displaying detailed information about all configured VPLS instances.

```
#show mpls vpls detail
Virtual Private LAN Service Instance: vpls3100, ID: 3100
SIG-Protocol: BGP
Route-Distinguisher :65010:3100
Route-Target :65010:3100
VE-ID :31
Attachment-Circuit :UP
Learning: Enabled
Group ID: 0, Configured MTU: 9216
Description: none
service-tpid: dot1.q
Operating mode: Raw
Configured interfaces:
Interface: xe26
Service-template : vpls3100_3100_13100
Match criteria : 3100
Action type : Translate
Action value : 4075
Outgoing tpid : dot1.q
Mesh Peers:
2.2.2.2 (Up)
```

The following table explains the show command output fields.

Table 77. show mpls vpls details output field

| Field | Description |
|--------------------------------------|--|
| Virtual Private LAN Service Instance | Number of VPLS service instance. |
| ID | VPLS identification detail for service instance. |
| SIG-Protocol | Type of protocol in MPLS configuration. |
| Attachment-Circuit | Details of the attached circuit in interface. |
| Learning | State of the interface. |
| Group ID | Group identification detail for VLAN. |

Table 77. show mpls vpls details output field (continued)

| Field | Description |
|-----------------------|---|
| VPLS Type | Type of VPLS in MPLS protocol. |
| Configured MTU | Number of configured MTU in the VPLs. |
| Description | Details of VPLS. |
| Service-tpid | Service TP identifier configured for the VPLS PW. |
| Operating mode | Type of mode in the interface. |
| Svlan Id | Configures a specific virtual LAN (VLAN). |
| Svlan Tpid | Service vlan TP identifier for the VPLS PW. |
| Redundancy admin role | Creating a Backup Administrator Role. |
| Redundancy oper role | Operational Role of the VPLS instance. |
| Configured interfaces | Details of the configured interfaces. |
| Interface | Selects an interface to configure. |
| Oper-state | Displays the current status of the cross-connect segment – Up or Down. |
| Service-template | Used to configure advanced service-related option. |
| Match criteria | Identifies prefix characteristics (network, BGP path attribute, nexthop, and so on) for a specific sequence. |
| Mesh Peers | Configuring the VPLS mesh peers. |
| PW Status Local | Used to perform limited local configuration changes, monitor device status and utilization, and simple local troubleshooting. |
| Remote | PW status of Remote end. |

The following is a sample output of the `show mpls vpls mesh` command displaying information about all the core VC connections for all VPLS instances.

```
#show mpls vpls mesh
VPLS-ID      Peer Addr      In-Intf      In-Label      Out-Intf      Out-Label      Lkps/St      PW-INDEX      SIG-
Protocol Status  Ecmp-Group
1
ve N/A      192.168.0.80   eth0         16            eth0         640          1/Up          1             BGP          Acti
1
ve N/A      192.168.0.90   eth1         18            eth1         642          1/Up          2             BGP          Acti
2
ve N/A      192.168.0.80   eth0         19            eth0         641          1/Up          1             BGP          Acti
2
ve N/A      192.168.0.90   eth1         17            eth1         643          1/Up          2             BGP          Acti
#
```

The following table explains the show command output fields.

Table 78. show mpls vpls mesh output field

| Field | Description |
|-----------|--|
| VPLS-ID | Identification details of the VPLS. |
| Peer Addr | IP address of the peer device. |
| In-Intf | Installed as a result of configuring an interface. |

Table 78. show mpls vpls mesh output field (continued)

| Field | Description |
|--------------|---|
| In-label | Displays the ingress (incoming interface) label for this segment. |
| Out-Label | Label received from downstream neighbor for route. |
| Network-Intf | Installed as a result of configuring an interface. |
| Lkps/St | Opcode and Status of the VPLS PW. |
| PW-INDEX | Pseudo wire index |
| SIG-Protocol | Signalling protocol |
| Status | Status of Pseudo wire |
| Ecmp-Group | Equal cost multi path group |

The following is a sample output of the `show mpls vpls spoke` displaying the spoke VC connection to the VPLS instance.

```
#show mpls vpls spoke
VPLS-ID   Virtual Circuit In-Intf   In-Label  Out-Intf   Out-Label Lkps/St
1         t100            eth2      20        eth2       640      1/Up
#
```

The following table explains the show command output fields.

Table 79. show mpls vpls spoke output field

| Field | Description |
|-----------------|---|
| VPLS-ID | Identification details of the VPLS. |
| Virtual Circuit | Used in transportation of data over a packet switch computer network. |
| In-Intf | Installed as a result of configuring an interface. |
| In-label | Displays the ingress (incoming interface) label for this segment. |
| Out-Label | Label received from downstream neighbor for route. |
| Network-Intf | Installed as a result of configuring an interface. |
| Lkps/St | Opcode and Status of the VPLS PW. |

The following is a sample output of `show mpls vpls count` displaying information about total, active and inactive vpls instances.

```
#show mpls vpls count
-----
Total VPLS instances   : 2
Active VPLS instances  : 2
Inactive VPLS instances : 0
-----
```

The following table explains the show command output fields.

Table 80. show mpls vpls count output field

| Field | Description |
|-------------------------|-----------------------------------|
| Total VPLS instances | Number of total VPLS instance. |
| Active VPLS instances | Number of active VPLS instance. |
| Inactive VPLS instances | Number of inactive VPLS instance. |

show mpls vpls mac-address

Use this command to display retrieved VPLS learning mac-addresses on MPLS enabled node.

Command Syntax

```
show mpls vpls mac-address (name NAME |) (interface IFNAME |) (peer A.B.C.D |) (count |)
```

Parameters

NAME

Specify the name of the vpls instance

count

Counts the number of MAC address learned

IFNAME

Specify the name of interface

A.B.C.D

Specify the peer address

Command Mode

Exec mode and Privileged Exec mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#show mpls vpls mac-address
VPLS-ID      MAC address      Learned from      Peer address
1            08:00:27:85:28:8a  eth1              1.1.1.1
1            08:00:27:99:91:1d  eth3              -
```

```
#show mpls vpls mac-address count
Total no of MAC addresses learnt :2
```

```
#show mpls vpls mac-address name vpls1
MAC address      Learned from      Peer address
08:00:27:85:28:8a  eth1              1.1.1.1
08:00:27:99:91:1d  eth3              -
```

```
#show mpls vpls mac-address name vpls1 count
Total no of MAC addresses learnt :2
```

```
#show mpls vpls mac-address interface eth1
VPLS-ID      MAC address      Learned from      Peer address
1            08:00:27:85:28:8a  eth1              1.1.1.1
```

```
#show mpls vpls mac-address interface eth1 count
Total no of MAC addresses learnt :1
```



```
#show mpls vpls mac-address name vpls1 interface eth1
MAC address      Learned from    Peer address
08:00:27:85:28:8a  eth1          1.1.1.1

#show mpls vpls mac-address name vpls1 interface eth1 count
Total no of MAC addresses learnt :1

#show mpls vpls mac-address peer 1.1.1.1
VPLS-ID      MAC address      Learned from    Peer address
1            08:00:27:85:28:8a  eth1          1.1.1.1

#show mpls vpls mac-address peer 1.1.1.1 count
Total no of MAC addresses learnt :1

#show mpls vpls mac-address name vpls1 peer 1.1.1.1
MAC address      Learned from    Peer address
08:00:27:85:28:8a  eth1          1.1.1.1

#show mpls vpls mac-address name vpls1 peer 1.1.1.1 count
Total no of MAC addresses learnt :1

#show mpls vpls mac-address interface eth1 peer 1.1.1.1
VPLS-ID      MAC address      Learned from    Peer address
1            08:00:27:85:28:8a  eth1          1.1.1.1

# show mpls vpls mac-address interface eth1 peer 1.1.1.1 count
Total no of MAC addresses learnt :1

#show mpls vpls mac-address name vpls1 interface eth1 peer 1.1.1.1
MAC address      Learned from    Peer address
08:00:27:85:28:8a  eth1          1.1.1.1

#show mpls vpls mac-address name vpls1 interface eth1 peer 1.1.1.1 count
Total no of MAC addresses learnt :1
```

The following table explains the show command output fields.

Table 81. show mpls vpls mac-address output field

| Field | Description |
|--------------|--|
| MAC address | Used to forward the packet into a given VPLS instance. |
| Learned from | MAC addresses learned from a specific interface. |
| Peer address | IP address of the peer device. |

show mpls vpls split-horizon

Use this command to display the split horizon grouping information.

Command Syntax

```
show mpls vpls split-horizon <vpls_name>
```

Parameters

<vpls_name>

(Optional) Specifies the name of the VPLS.

Default

None

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 6.5.2.

Example

The following example shows a summary of all VPLS instances, showing their VPLS ID, type, details (such as spokes or ACs), associated interfaces, and the split horizon group they belong to:

```
#show mpls vpls split-horizon
+-----+-----+-----+-----+-----+
| VPLSID | Type | Details | Interface | Group |
+-----+-----+-----+-----+-----+
800      AC      st2      xe9         access2
800      AC      -        xe17.800   access1
800      Spoke   vc2      -          network
2000     Spoke   vc1      -          access1
```

The following example shows the split horizon details for a specific VPLS instance, showing similar information but focused on the specified VPLS ID:

```
#show mpls vpls split-horizon vpls800
+-----+-----+-----+-----+-----+
| VPLSID | Type | Details | Interface | Group |
+-----+-----+-----+-----+-----+
800      AC      st2      xe9         access2
800      AC      -        xe17.800   access1
800      Spoke   vc2      -          network
```

The following example shows when there is no VPLS is configured with the specified VPLS ID:

```
#show mpls vpls split-horizon vpls200
% No VPLS with name: vpls200 configured.
```

The following example shows VPLS split horizon details for VPLS ID 2000:

```
#show mpls vpls split-horizon vpls2000
```

| VPLSID | Type | Details | Interface | Group |
|--------|-------|---------|-----------|---------|
| 2000 | Spoke | vc1 | - | access1 |

The following table explains the show command output fields.

Table 82. show mpls vpls split-horizon output field

| Field | Description |
|-----------|--|
| VPLSID | The ID of the VPLS instance. It uniquely identifies each VPLS in the network. |
| Type | The type of connection within the VPLS instance. |
| Details | Additional details about the VPLS connection type, such as the virtual circuit (VC) number or other identifiers related to the connection. |
| Interface | The physical or logical interface on the PE router that is part of the VPLS instance. |
| Group | The split-horizon group associated with the VPLS instance. |

show mpls vpls statistics

Use this command to display MPLS traffic statistics for VPLS network or access or all ports.



Note: Multicast traffic statistics not supported by hardware.

Command Syntax

```
show mpls vpls NAME statistics
show mpls vpls NAME statistics ((network-port ((peer A.B.C.D)|(spoke-vc VC-NAME)|)) | (access-port
(IFNAME (ethernet|(vlan <1-4094>)))))
```

Parameters

NAME

Name of the VPLS instance

a.b.c.d

Mesh peer address of VC instance

VC-NAME

Name of the spoke VC instance

IFNAME

Name of the access-port interface

<1-4094>

VLAN ID of access-port of type VLAN

Command Mode

Execution mode and Privileged execution mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#show mpls vpls v1 statistics
Virtual Private LAN Service Instance: v1, ID: 10

Access port statistics:
Interface: xe3/4 VLAN ID: 2
  RX:  Input packets  : 10
      Input bytes    : 640
  TX:  Output packets : 0
      Output bytes   : 0

Network port statistics:
Mesh Peer: 8.8.8.8 (Up)
  RX:  Input packets  : 0
      Input bytes    : 0
  TX:  Output packets : 10
      Output bytes   : 640
```

The following table explains the show command output fields.

Table 83. show mpls vpls statistics output field

| Field | Description |
|-------------------------|---|
| Access port statistics | Traffic statistics on Access port of VC/VPLS. |
| Network port statistics | Traffic statistics on Provider port of VC/VPLS. |
| Interface | Type of interface in the network. |
| VLAN ID | Identification details of the VPLS. |
| Mesh Peer | Configuring the VPLS mesh peers. |
| RX | Number of received packets. |
| Input packets | Number of hello packets received from the neighbor. |
| Input bytes | Size of hello packets received from the neighbor. |
| TX | Number of packets transmitted. |
| Output packets | Number of hello packets sent to the neighbor. |
| Output bytes | Size of hello packets sent to the neighbor. |

signaling ldp

Use this command to establish a pseudowire connection between Provider Edge (PE) routers. Use this command to use the Label Distribution Protocol (LDP) for signaling and to support VPLS auto-discovery between VPLS instances. Using this command triggers LDP to signal a pseudowire between the configured VPLS peers in the same VPLS instance. The `vpls-peer` command is used to identify the VPLS peers that are part of a VPLS instance.



Note: Issuing this command puts the router into VPLS signaling (`config-vpls-sig`) mode.

Use the `no` parameter with this command to remove (tear down) pseudowires with other PE routers.

Command Syntax

```
signaling ldp
no signaling ldp
```

Parameters

None

Default

By default, signaling ldp is disabled.

Command Mode

VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)# mpls vpls test 100
(config-vpls)#signaling ldp
(config-vpls-sig)#vpls-peer 97.97.97.97
(config-vpls-sig)#exit
```

signaling bgp

Use this command to establish a pseudowire connection between Provider Edge (PE) routers. Use this command to use the Border Gateway Protocol (BGP) for signaling and to support VPLS auto-discovery between VPLS instances. Using this command triggers BGP to auto-discover VPLS peers and signal pseudowire between the VPLS peers in the same VPLS instance.



Note: Issuing this command puts the router into VPLS signaling .

Use the `no` parameter with this command to remove (tear down) pseudowires with other PE routers.

Command Syntax

```
signaling bgp
no signaling bgp
```

Parameters

None

Default

By default, signaling bgp is disabled.

Command Mode

VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)# mpls vpls test 100
(config-vpls)#signaling bgp
(config-vpls-sig)#exit
```

static-mac

Use this command to add static MAC address to attachment circuit specific for a VPLS instance.

Use the `no` parameter with this command to remove static MAC address.



Note: It is not supported, if the user configures same mac address on different attachment circuits for same VPLS instance.

Command Syntax

```
static-mac XXXX.XXXX.XXXX  
no static-mac XXXX.XXXX.XXXX
```

Parameter

XXXX.XXXX.XXXX

MAC address in HHHH.HHHH.HHHH format.

Default

By default, mac is disabled.

Command Mode

Interface VPLS

Applicability

This command was introduced before OcNOS version 4.2.

Examples

```
(config)#interface cell/2  
(config-if)#mpls-vpls vpls2 service-template vc1  
(config-if-vpls)#static-mac 0000.0400.0602
```

tunnel-select-policy

Use this command to set tunnel-policy for a VPLS instance. This command is used for BGP signaling based VPLS instance. This will be used while selecting transport for the peer.

Use the `no` parameter with this command to un-configure tunnel selection policy.

Command Syntax

```
tunnel-select-policy TNLPOLICYNAME  
no tunnel-select-policy
```

Parameters

TNLPOLICYNAME

Name of tunnel policy

Command Mode

VPLS signal mode

Applicability

This command was introduced before OcNOS version 4.2.

Examples

```
#configure terminal  
(config)#mpls vpls vpls1 10  
(config-vpls)#signaling bgp  
(config-vpls-sig)#tunnel-select-policy policy1
```

ve-id

Use this command to configure a VPLS Edge (VE) device. Each Provider Edge (PE) device participating in a VPLS must have at least one VE ID. When the PE is connected to several u-PEs (Layer 2 PE devices used to provide Layer 2 aggregation), there are unique VE ID's for each u-PE. The PE may also be assigned a VE ID, if it is to act as the VE for the VPLS.

Use the `no` parameter with this command to remove a VE ID.



Notes:

- The VE-range is not configurable and set to a fixed value of 64 where VE refers to the VPLS edge device.
- The VE-IDs configured at the VPLS peer nodes must be a positive number and must always lie in a range of plus or minus 64 (VE-range) of the local VE-ID.

Command Syntax

```
ve-id <1-65535>  
no ve-id <1-65535>
```

Parameters

<1-65535>

VE-ID's range is between 1 and 65535. This should be unique among the VPLS Peers for a VPLS instance.

Default

By default, ve id is disabled.

Command Mode

BGP VPLS Signaling mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#mpls vpls test 100  
(config-vpls)#signaling bgp  
(config-vpls-sig)#ve-id 2  
(config-vpls-sig)#exit
```

vpls-ac-group

Use this command to assign an Attachment Circuit (AC) group to VPLS.

Use the `no` parameter with this command to remove an AC group.

Command Syntax

```
vpls-ac-group GROUPNAME  
no vpls-ac-group
```

Parameter

GROUPNAME

Enter a name for the AC group

Default

By default, vpls ac group is disabled.

Command Mode

VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#mpls vpls test 12  
(config-vpls)#vpls-ac-group new-ac  
(config-vpls)#no vpls-ac-group
```

vpls-description

Use this command to add a description line for a VPLS instance.

Use the `no` parameter with this command to remove a VPLS description.

Command Syntax

```
vpls-description LINE  
no vpls-description (LINE|)
```

Parameter

LINE

Enter a text string for the VPLS instance

Default

By default, vpls description is disabled.

Command Mode

VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#mpls vpls test 34  
(config-vpls)#vpls-description This is for testing  
(config-vpls)#exit
```

vpls fib-entry

Use this command to create a static VPLS FIB entry. When a VPLS peer is configured manually, no signaling is done. Therefore, a VPLS static entry must be created for all manually created nodes.

Use the `no` option with this command to delete a static VPLS FIB entry.

Command Syntax

```
vpls fib-entry VPLS-ID (peer A.B.C.D| spoke-vc VC-NAME) IN-LABEL OUT-INTF OUT-LABEL
no vpls fib-entry VPLS-ID ((peer A.B.C.D| (spoke-vc VC-NAME))
no vpls fib-entry VPLS-ID ((peer A.B.C.D| (spoke-vc VC-NAME)) IN-LABEL OUT-INTF OUT-LABEL
```

Parameters

VPLS-ID

VPLS identifier

peer

Mesh peer address VPLS identifier

A.B.C.D

Peer IPv4 Address.

spoke-vc

Spoke VC

VC-NAME

Virtual Circuit name

IN-LABEL

Incoming label value in the range of <16-15999>

OUT-INTF

Provider-facing interface

OUT-LABEL

Outgoing label value in the range of <16-15999>

Default

By default, vpls fib entry is disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The first example shows how to configure VPLS FIB entry 100 with mesh peer 97.97.97.97 for incoming label 15999, outgoing interface eth2 with outgoing label 15999:

```
#configure terminal
(config)#vpls fib-entry 100 peer 97.97.97.97 15999 eth2 15999
```

The second example shows how to configure VPLS FIB entry 100 with spoke-vc t1 for incoming label 15999, outgoing interface eth2 with outgoing label 15999:

```
#configure terminal
(config)#vpls fib-entry 100 spoke-vc t1 15999 eth2 15999
```

vpls-mtu

Use this command to set the Maximum Transmission Unit (MTU) size for a given VPLS instance. This size is signaled to peer VPLS routers.

Use the `no` parameter with this command to remove the MTU size setting.

Command Syntax

```
vpls-mtu <576-65535>  
no vpls-mtu (<576-65535>|)
```

Parameter

<576-65535>

Range of MTU size allowed for a VPLS instance

Default

By default, vpls mtu is 1500.

Command Mode

VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#mpls vpls test 34  
(config-vpls)#vpls-mtu 6506  
(config-vpls)#exit
```

vpls-peer

Use this command to add a peer to a VPLS domain. This command triggers Label Distribution Protocol (LDP) signaling by default.

Use the **no** parameter to delete a VPLS virtual circuit for a specific peer.

Use the **no** parameter to delete a VPLS tunnel mapping for a specific peer to delete/unmap tunnel-id, tunnel-name or tunnel-select-policy for a vpls-peer.

Command Syntax

```
vpls-peer A.B.C.D ((agi NAME saii NAME taii NAME) |) ((tunnel-id <1-5000>) | (tunnel-name TNLNAME) |)
((tunnel-select-policy TNLPOLICYNAME) |)
no vpls-peer A.B.C.D ((tunnel-id | tunnel-name | tunnel-select-policy) |)
```

Parameters

A.B.C.D

The address of a VPLS peer node to which a mesh virtual circuit is to be created

tunnel-id

Static Tunnel Identifier

<1-5000>

Identifying value for Tunnel-id

A.B.C.D

IPv4 Address for end-point for FEC129 MPLS Layer-2 Virtual Circuit

agi

Specify the value used for the AGI in FEC129 MPLS Layer-2 Virtual Circuit

NAME

AGI value for FEC129 MPLS Layer-2 Virtual Circuit

saii

Specify the value used for the SAIL in FEC129 MPLS Layer-2 Virtual Circuit

NAME

SAIL value for FEC129 MPLS Layer-2 Virtual Circuit

taii

Specify the value used for the TAIL in FEC129 MPLS Layer-2 Virtual Circuit

NAME

TAIL value for FEC129 MPLS Layer-2 Virtual Circuit

tunnel-name

Tunnel name of the MPLS LSP (or Layer 2 Tunnel) to be used

TUNNELNAME

Identifying name for Tunnel

tunnel-select-policy

Tunnel selection policy

TNLPOLICYNAME

Name of tunnel policy

Default

By default, vpls peer is disabled.

Command Mode

VPLS signal mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#mpls vpls test 100
(config-vpls)#signaling ldp
(config-vpls-sig)#vpls-peer 97.97.97.97
(config-vpls-sig)#vpls-peer 97.97.97.97 tunnel-id 24
(config-vpls)#exit
(config)#exit
```

vpls-peer manual

Use this command to statically configure a VPLS peer. Because this command is not used in signaling mode, no signaling is used to set up the virtual circuit. At least one such peer configuration is required for every VPLS instance.

Use the `no` parameter with this command to remove a statically configured VPLS peer.

Use the `no` parameter to delete a VPLS tunnel mapping for a specific peer to delete/unmap tunnel-id, tunnel-name or tunnel-select-policy for a vpls-peer.

Command Syntax

```
vpls-peer A.B.C.D ((tunnel-id <1-5000>) | (tunnel-name TNLNAME) |) (manual|)  
no vpls-peer A.B.C.D ((tunnel-id | tunnel-name | tunnel-select-policy) |)
```

Parameters

A.B.C.D

The address of a VPLS peer node to which a mesh virtual circuit is to be created

tunnel-id

Static Tunnel Identifier

<1-5000>

Identifying value for Tunnel-id

tunnel-name

Tunnel name of the MPLS LSP (or Layer 2 Tunnel) to be used

TUNNELNAME

Identifying name for Tunnel

Default

By default, vpls peer A.B.C.D manual is disabled

Command Mode

VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal  
(config)#mpls vpls test 100  
(config-vpls)#vpls-peer 97.97.97.97 manual  
(config-vpls)#vpls-peer 97.97.97.97 tunnel-id 24 manual  
(config-vpls)#exit  
(config)#exit
```

vpls-type

Use this command to assign a type (either Ethernet or VLAN) for VPLS.



Note: The default type is chosen as Ethernet.

Command Syntax

```
vpls-type (ethernet|vlan)
```

Parameter

ethernet

Designate Ethernet as the VPLS type

vlan

Designate VLAN as the VPLS type

Default

By default, vpls type is ethernet.

Command Mode

signaling ldp mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#mpls vpls test 100
(config-vpls)#signaling ldp
(config-vpls-sig)#vpls-type vlan
(config-vpls-sig)#vpls-peer 2.2.2.2
(config-vpls-sig)#exit
(config-vpls)#exit
```

vpls-vc

Use this command add a spoke virtual circuit to VPLS domain.

Use the `no` parameter to remove this configuration.

Command Syntax

```
vpls-vc NAME (ethernet|vlan|)  
vpls-vc NAME (secondary NAME|) (ethernet|vlan|)  
no vpls-vc NAME
```

Parameter

NAME

Enter a string that identifies the MPLS VC to add to the VPLS domain

secondary

Set the secondary spoke name

NAME

Enter a string that identifies the secondary spoke

ethernet

Identify the spoke type as Ethernet (default)

vlan

Identify the spoke type as VLAN.

TNLNAME

Specify the MPLS-TP tunnel-name.

Default

By default, vpls vc name is disabled.

Command Mode

VPLS mode

Applicability

This command was introduced before OcNOS version 1.3.

Example

```
#configure terminal  
(config)#mpls vpls test 34  
(config-vpls)#vpls-vc VC1  
(config-vpls)#exit  
(config)#exit
```

vpls-route-map

Use this command to configure the route-map name on an vpls instance.

Use the `no` parameter to remove the route-map configured.



Note: This command imports the extended-community color value configured on the route-map.

This command is used for BGP signaling based VPLS instance and not applicable for LDP signaling.

Command Syntax

```
vpls-route-map WORD  
no vpls-route-map
```

Parameter

WORD

Name of the route-map.

Default

None

Command Mode

VPLS mode

Applicability

This command was introduced before OcNOS version 6.6.0.

Example

```
OcNOS(config)#mpls vpls a 1  
OcNOS(config-vpls)#vpls-route-map map1  
OcNOS(config-vpls)#exit  
OcNOS(config)#  
  
OcNOS(config)#mpls vpls a 1  
OcNOS(config-vpls)#no vpls-route-map  
OcNOS(config-vpls)#
```

EVPN MPLS Commands

This chapter describes the EVPN MPLS commands:

| | |
|--|------|
| access-if-evpn | 1540 |
| arp-cache disable | 1541 |
| arp-nd cos | 1542 |
| arp-nd flood-suppress | 1543 |
| arp-nd refresh timer | 1544 |
| clear evpn mpls counters | 1545 |
| clear mac address-table | 1546 |
| dynamic-learning disable | 1547 |
| evi-name | 1548 |
| evpn esi hold-time | 1549 |
| evpn etree | 1550 |
| evpn irb | 1551 |
| evpn irb-forwarding anycast-gateway-mac | 1552 |
| evpn irb-if-forwarding anycast-gateway-mac | 1553 |
| evpn mpls enable | 1554 |
| evpn mpls id | 1555 |
| evpn mpls irb | 1557 |
| l3vni | 1558 |
| mac | 1559 |
| evpn mpls mac-ageing-time | 1560 |
| evpn mpls multihoming enable | 1561 |
| evpn mpls vtep-ip-global | 1562 |
| evpn multi-homed | 1563 |
| hardware-profile filter evpn-mpls-cw | 1565 |
| hardware-profile filter evpn-mpls-mh | 1566 |
| hardware-profile statistics evpn-mpls | 1567 |
| host-reachability-protocol | 1568 |
| evpn-route-map | 1569 |
| ip address | 1570 |
| ipv6 address | 1571 |
| link-loss forwarding | 1572 |
| mac-holdtime | 1573 |
| mac-vrf | 1574 |
| map vpn-id | 1575 |
| nd-cache disable | 1576 |
| rewrite egr-pcp-disable | 1577 |

| | |
|--|------|
| service-carving | 1578 |
| service-carving ac-driven | 1580 |
| service-carving weight | 1581 |
| show bgp l2vpn evpn | 1582 |
| show evpn load-balance | 1584 |
| show evpn mpls | 1586 |
| show evpn mpls arp-cache | 1587 |
| show evpn mpls counters | 1588 |
| show evpn mpls label alias | 1589 |
| show evpn mpls label esi | 1590 |
| show evpn mpls mac-table | 1591 |
| show evpn mpls nd-cache | 1593 |
| show evpn mpls route-count | 1594 |
| show evpn mpls static host state | 1596 |
| show evpn mpls status | 1597 |
| show evpn mpls tunnel | 1598 |
| show evpn mpls tunnel label | 1599 |
| show evpn mpls xconnect | 1600 |
| show evpn mpls xconnect tunnel | 1601 |
| show evpn mpls xconnect tunnel label | 1602 |
| show running-config evpn mpls | 1603 |
| tunnel-select-policy | 1605 |

access-if-evpn

Use this command to create the evpn mpls access-port.

Use the `no` form of this command to delete the evpn mpls access-port.

Command Syntax

```
access-if-evpn  
no access-if-evpn
```

Parameters

None

Command Mode

L2 Sub-interface mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#configure terminal  
(config)#interface xel.1 switchport  
(config-if)#access-if-evpn  
(config-access-if)#end
```


arp-cache disable

Use this command to disable the ARP cache for MAC/IP.

When the ARP cache is disabled on a access circuit, It does not reply to any ARP arriving on this port from the cache. It withdraws all MAC/IPs configured/learned on this access circuit and removes the MAC/IP entry for this AC from the local ARP cache.

It also makes sure that on withdrawing the MAC/IP route, the MAC does not become unknown. If all routes for this MAC are being withdrawn because of this command, then It advertises a MAC-only route. This is done so that the MAC does not become unknown and only the cache functionality becomes disabled.

Use the `no` form of this command to enable ARP cache for MAC/IP.



Notes:

- On enabling the cache, an IP will be in conflict, then the cache enable will fail. The conflict has to be manually removed and then the cache enabled.
- When encapsulation default is configured under L2 subifp , then arp-packets will not be uplifted.
- Not applicable for the AC port which is mapped with ELINE/Xconnect Service.

Command Syntax

```
arp-cache disable
no arp-cache disable
```

Parameters

None

Default

By default, the arp-cache option is enabled.

Command Mode

Access if evpn mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
#configure terminal
(config)#interface xe1.1 switchport
(config-if)#access-if-evpn
(config-access-if)#arp-cache disable
(config-access-if)#end
```

arp-nd cos

Use this command to explicitly set or override the COS value (0–7) in the ARP/ND reply sent (also used in ARP/ND proxy sub-feature). By default ARP/ND reply will use CoS value of received ARP/ND request.

Use the no form of this command to use the default COS value.

Command Syntax

```
arp-nd cos <0-7>  
no arp-nd cos
```

Parameters

<0-7>

Default

CoS 0 or previous ARP/ND reply CoS being used.

Command Mode

Access if evpn mode

Applicability

This command was introduced in OcNOS version 6.5.3.

Examples

```
#configure terminal  
(config)#interface xel.1 switchport  
(config-if)#access-if-evpn  
(config-access-if)#arp-nd cos 5  
(config-access-if)#end
```

arp-nd flood-suppress

Use this command to completely restrict the flood of ARP/ND packets towards remote PEs or other Access Circuit

This command applies only when the ARP cache and ND cache are enabled. When the ARP cache is disabled, ARP flooding is not suppressed even if this command is given. When the ND cache is disabled, ND flooding is not disabled, even if this command is given.

Use the `no` form of this command to not restrict the flood of ARP/ND packets.



Note: Not applicable for the AC port which is mapped with ELINE/Xconnect Service.

Command Syntax

```
arp-nd flood-suppress
no arp-nd flood-suppress
```

Parameters

None

Default

By default, the arp-nd flood-suppress option is disabled.

Command Mode

Access if evpn mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
#configure terminal
(config)#interface xel.1 switchport
(config-if)#access-if-evpn
(config-access-if)#arp-nd flood-suppress
(config-access-if)#end
```

arp-nd refresh timer

Use this command to configure aging out the arp-cache and nd-cache entries for given time multiplied by 3 in seconds.

Use the `no` form of this command to remove the configuration.



Notes:

- Not applicable for the AC port which is mapped with ELINE/Xconnect Service.
- After this timer interval, it sends out ARP to revalidate and 3 times of this would lead to removal of the dynamic entry.

Command Syntax

```
evpn mpls arp-nd refresh-timer <3-190>  
no evpn mpls arp-nd refresh-timer
```

Parameters

<3-190>

Refresh timer value in seconds (age-out is refresh time * 3)

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
#config mode  
(config)#evpn mpls arp-nd refresh-timer 100  
(config)#no evpn mpls arp-nd refresh-timer
```

clear evpn mpls counters



Note: Use this command to clear the counters of a network port.

Command Syntax

```
clear evpn mpls counters ((network (ingress| (egress dst <A.B.C.D>))) | vpn-id <1-16777215> (network (ingress| (egress dst <A.B.C.D>))))
```

Parameters

<1-16777215>

VPN identifier

A.B.C.D

Ipv4 address

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
DUT2#clear evpn mpls counters network egress dst 105.1.1.1
DUT2#
DUT2#clear evpn mpls counters vpn-id 10 network ingress
DUT2#
```

clear mac address-table

Use this command to clear dynamically learned MACs.



Note: To make evpn mpls disable and enable effective, system reboot is required.

Command Syntax

```
clear mac address-table dynamic evpn-mpls
clear mac address-table dynamic evpn-mpls evid <1-16777215>
clear mac address-table dynamic evpn-mpls evid <1-16777215> (address MACADDR|)
```

Parameters

address

Clear the specified MAC Address

evid

EVPN-MPLS tenant identifier

<1-16777215>

Range supported for EVID

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
#clear mac address-table dynamic evpn-mpls
#clear mac address-table dynamic evpn-mpls evid 30
OcNOS#clear mac address-table dynamic evpn-mpls evid 30 address 0000.0022.2222
```

dynamic-learning disable

Use this command to disable dynamic learning of MACs at the Access Circuit.

This command also disables dynamic learning of MAC/IP from ARP/ND messages received on this Access Circuit.

Use the `no` form of this command to enable dynamic learning of MACs at the Access Circuit.



Notes:

- Not applicable for the AC port which is mapped with ELINE/Xconnect Service.
- IPv4 and IPv6 addresses associated with dynamic MACs also get cleaned with this command.

Command Syntax

```
dynamic-learning disable  
no dynamic-learning disable
```

Parameters

None

Default

By default, the dynamic-learning option is enabled.

Command Mode

Access if evpn mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#configure terminal  
(config)#interface xel.1 switchport  
(config-if)#access-if-evpn  
(config-access-if)#dynamic-learning disable  
(config-access-if)#end
```

evi-name

Use this command to name the EVPN MPLS ID.

Use the `no` form of this command to remove the name of the EVPN MPLS ID.

Command Syntax

```
evi-name <WORD>  
no evi-name
```

Parameters

WORD

EVI name. Maximum limit 10 characters (shall not be only numeric).

Default

None

Command Mode

EVPN MPLS mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#configure terminal  
(config)#evpn mpls id 3  
(config-evpn-mpls)#evi-name vni_ELAN  
(config-evpn-mpls)#exit
```

evpn esi hold-time

Use this command to allow some time for the evpn mpls tunnels and xconnect tunnel to come at the time of evpn mpls initialization before making the ESI up. This avoids traffic to be black-holed when a new PE is added and connected to an already running CE for multi-homing.

Use the `no` form of this command to make the esi up immediately when configuring the `access-if` cli.

Command Syntax

```
evpn esi hold-time <10-300>
no evpn esi hold-time
```

Parameters

<10-300>

Hold time in seconds

Default

Default value is 0.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#configure terminal
(config)# evpn esi hold-time 100
(config)# exit
```

evpn etree

Use this command to enable E-Tree functionality within the EVPN configuration. The `evpn etree enable` option ([EVPN MPLS E-Tree Scenario 1 \(page 1859\)](#)) configures the device to function as either a Leaf or a Root site per EVPN Instance (EVI), as defined in RFC-8317. The `scenario-2` option ([EVPN MPLS E-Tree Scenario 2 \(page 1888\)](#)) configures the device to support both Root and Leaf sites on the same Provider Edge (PE) device, enforcing traffic filtering rules to prevent Leaf-to-Leaf communication while allowing Root-to-Root and Root-to-Leaf traffic.



Note: If Scenario 1 is configured, it must be removed before enabling Scenario 2.

Command Syntax

```
evpn etree enable (scenario-2|)
no evpn etree enable (scenario-2|)
```

Parameters

scenario-2

Enables EVPN MPLS E-Tree Scenario 2, which configures the device to support both Root and Leaf sites on the same PE device.

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.1. Introduced the `scenario-2` parameter in OcNOS version 7.0.0.

Example

The following example illustrates how to activate E-Tree Scenario 1 functionality for EVPN, which configures the device as either a Leaf or a Root site per EVI.

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable

OcNOS(config)#no evpn etree enable
```

The following example illustrates how to activate E-Tree Scenario 2 functionality for EVPN.

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable scenario-2
```

evpn irb

Use this command to map an IRB interface to an L2 VNID. This IRB interface can have multiple IP address as explained in IRB IP address CLI and can serve all subnets attached to the L2 VNID. Currently its 1:1 mapping between IRB interface & L2VNI.

This configuration enables default gateway behavior on that VTEP for that VNID and subnet and triggers default gateway MAC-IP route generation for the corresponding IRB IP with the local or global IRB MAC (router MAC) in distributed gateway in distributed gateway

This makes the default behavior as enable for gateway functionality.

No command disables default gateway behavior of the VTEP for that VNID subnet and triggers withdrawal of the default gateway MAC/IP advertisement done for that VNID subnet.

Use the `no` form this command to disable default gateway behavior of the VTEP for that VNID subnet and triggers withdrawal of the default gateway MAC/IP advertisement done for that VNID subnet.

Command Syntax

```
evpn irb [<NAME>|irb-advertise-host-route]
no evpn irb [<NAME>|irb-advertise-host-route]
```

Parameters

NAME

IRB interface name

irb-advertise-host-route

To advertise host IP prefixes

Command Mode

EVPN_MPLS_CONFIG Mode

Applicability

This command was introduced before OcNOS version 6.0.0.

irb-advertise-host-route introduced in OcNOS version 6.2.0.

Example

```
#config mode
(config)#evpn mpls id 10
(config-evpn-mpls)#evpn irb irb100

(config)#evpn mpls id 102
(config-evpn-mpls)#evpn irb-advertise-host-route
```

evpn irb-forwarding anycast-gateway-mac

Use this command to configure common anycast mac-address for all the IRB interfaces

Use the `no` form of this command to remove the global MAC address on all the IRB interfaces.

Command Syntax

```
evpn irb-forwarding anycast-gateway-mac MAC
no evpn irb-forwarding anycast-gateway-mac
```

Parameters

MAC

XX-XX-XX-XX-XX-XX Source MAC address (Option 1)

XX:XX:XX:XX:XX:XX Source MAC address (Option 2)

XXXX.XXXX.XXXX Source MAC address (Option 3)

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 4.1.

Example

```
#configure terminal
(config)#evpn irb-forwarding anycast-gateway-mac 0000.0000.1313
```

Or

```
(config)#evpn irb-forwarding anycast-gateway-mac 00:00:00:00:13:13
```

Or

```
(config)#evpn irb-forwarding anycast-gateway-mac 00-00-00-00-13-13
(config)#no evpn irb-forwarding anycast-gateway-mac
```

evpn irb-if-forwarding anycast-gateway-mac

Use this command to enable a Layer 3 interface to use the global anycast IRB MAC address.

Use the `no` form of this command to unconfigure the anycast MAC at a Layer 3 interface.

Command Syntax

```
evpn irb-if-forwarding anycast-gateway-mac
no evpn irb-if-forwarding anycast-gateway-mac
```

Parameters

None

Command Mode

IRB interface mode

L3 interface mode

Applicability

This command was introduced before OcNOS version 6.0.0.

L3 interface mode support was added in OcNOS version 6.1.0.

Example

```
#configure terminal
(config)#interface irb 1
(config-irb-if)#ip vrf forwarding vrfip
(config-irb-if)#evpn irb-if-forwarding anycast-gateway-mac
(config-irb-if)#no evpn irb-if-forwarding anycast-gateway-mac
(config)#interface xel
(config-irb-if)#ip vrf forwarding vrfip
(config-irb-if)#evpn irb-if-forwarding anycast-gateway-mac
(config-irb-if)#no evpn irb-if-forwarding anycast-gateway-mac
```

evpn mpls enable

Use this command to enable evpn mpls functionality.

Use the no version of this command to disable evpn mpls functionality.



Note: To make evpn mpls disable and enable effective, system reboot is required.

Command Syntax

```
evpn mpls enable  
no evpn mpls enable
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
(config)#evpn mpls enable  
(config)#no evpn mpls enable
```

evpn mpls id

Use this command to set the VPN identifier to create an EVPN MPLS tunnel.



Note: To set the VPN identifier for an E-LAN, use the `evpn mpls id` command. For creating an E-LINE/XConnect, use the `evpn mpls xconnect` command with source and target identifiers.

Use the `no` parameter of this command to delete the EVPN MPLS ID for the MPLS tunnel.

Command Syntax

```
evpn mpls id <1-16777215> (| xconnect target-mpls-id <1-16777215>) (|control-word) (|etree-leaf)
(|flow-label)
no evpn mpls id <1-16777215>
```

Parameters

evpn mpls id <1-16777215>

Specifies the EVPN-MPLS tenant identifier. This is a numeric value ranging from 1 to 16777215.

xconnect target-mpls-id <1-16777215>

Enables E-LINE Xconnect. Specifies the target EVID for E-LINE Xconnect. This is a numeric value within the range from 1 to 16777215.

control-word

Enables control-word egress or ingress options for the given EVPN Instance (E-LAN or E-LINE).

etree-leaf

(Optional) Configures the device as a leaf node within the E-Tree topology.

flow-label

Inserts a [Flow-Aware Transport \(FAT\) label](#) for per-flow hashing to improve ECMP traffic distribution.

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 3.0.

Introduced `xconnect target-mpls-id <1-16777215>` parameter in the OcNOS version 4.0, `control-word` parameter in the OcNOS version 6.0.0, `etree-leaf` parameter in the OcNOS version 6.5.1, and `flow-label` parameter in the OcNOS version 7.0.0.

Example

```
(config)#evpn mpls id 10
(config-evpn-mpls)#exit

(config)#no evpn mpls id 10

(config)#evpn mpls id 100 xconnect target-mpls-id 200
(config-evpn-mpls)#exit
```

```
(config)#no evpn mpls id 100

(config)#evpn mpls id 300 xconnect target-mpls-id 400 control-word
(config-evpn-mpls)#exit

(config)#no evpn mpls id 300
```

Use the following command to configure the leaf node as an E-Tree leaf in a MPLS EVPN network.

```
(config)#evpn mpls id 10 etree-leaf
(config-evpn-mpls)#exit
```

The following configuration shows how to enable the `flow-label` parameter under different EVPN MPLS ID instances:

```
!
evpn mpls id 10 xconnect target-mpls-id 15 flow-label
  host-reachability-protocol evpn-bgp vrf2
!
evpn mpls id 20 xconnect target-mpls-id 25 control-word flow-label
  host-reachability-protocol evpn-bgp vrf2
!
evpn mpls id 100 flow-label
  host-reachability-protocol evpn-bgp vrf1
!
evpn mpls id 200 control-word flow-label
  host-reachability-protocol evpn-bgp vrf1
!
```

evpn mpls irb

Use this command to enable EVPN-MPLS-IRB globally. Applicable to EVPN-MPLS

Use the `no` form this command to disable EVPN-MPLS.

Command Syntax

```
evpn mpls irb
no evpn mpls irb
```

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 6.0.0.

Example

```
#config mode
(config)#evpn mpls irb
(config)#no evpn mpls irb
```

l3vni

Use this command to configure L3 Virtual Network Identifier for an ip vrf

Use the no form of this command to remove L3 Virtual Network Identifier

This identifies a tenant, with this one tenant can have L3VNI as its identifier and he can have multiple L2 networks identified with L2VNI's.



Note: L3 VNID cannot be same as L2 VNID.

Command Syntax

```
l3vni <L3 VNID>  
no l3vni <L3 VNID>
```

Parameters

<1-16777215>

L3 VNID. Cannot be same as L2 VNID

Command Mode

Configure VRF mode

Applicability

This command was introduced in OcNOS version 4.1.

Example

```
#configure terminal  
(config)#ip vrf vrfip  
(config-vrf)#l3vni 10002  
(config-vrf)#no l3vni 10002
```

mac

Use this command to configure a static MAC address with IPv4/IPv6 address or only MAC address under the Access-Circuit.

Use the no form of this command to unconfigure a static MAC address with IPv4/IPv6 address and only MAC address from the Access-Circuit.



Note: Not applicable for the AC port which is mapped with ELINE Service.

For static mac addresses,

- 1) the local static is given preference over remote learned static.
- 2) if same static mac is configured on more than one interface of same VPN then conflict state need to be resolved manually by removing that static-mac on all access-interfaces where it was configured.

Command Syntax

```
mac XXXX.XXXX.XXXX (| ip A.B.C.D | ipv6 X:X::X:X)
no mac XXXX.XXXX.XXXX (| ip A.B.C.D | ipv6 X:X::X:X)
```

Parameters

XXXX.XXXX.XXXX

Static mac address

A.B.C.D

Static IPv4 address

X:X::X:X

Static IPv6 address

Command Mode

Access if evpn mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
#configure terminal
(config)# interface xel.1 switchport
(config-if)#access-if-evpn
(config-access-if)#mac 0000.0000.1001
(config-access-if)#mac 0000.0000.1002 ip 10.10.10.1
(config-access-if)#mac 0000.0000.1003 ipv6 1201::1
(config-access-if)#end
```

evpn mpls mac-ageing-time

Use this command to set the dynamically learned MAC aging time.

Use the `no` form of this command to set the age out the MACs in hardware to its default (300 seconds).



Note: This command affects the default bridge ageing time.

Command Syntax

```
evpn mpls mac-ageing-time <10-572>  
no evpn mpls mac-ageing-time
```

Parameters

<10-572>

Ageing time in seconds.

Default

The default age out time is 300 seconds.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#configure terminal  
(config)#evpn mpls mac-ageing-time 10
```

evpn mpls multihoming enable

Use this command to enable evpn mpls multihoming

Use the `no` form of this command to disable evpn mpls multihoming.



Notes:

- Node will have to be restarted for this to be applicable. If there are some nodes in topology which have multi-homed CEs, then the nodes which do not have multi-homed CEs should also enable multihoming, so that they can load share traffic to the multi-homed CEs.
- Before enabling multi-homing, configure the hardware-profiles with the [hardware-profile filter evpn-mpls-mh \(page 1566\)](#) commands.

Command Syntax

```
evpn mpls multihoming enable
no evpn mpls multihoming enable
```

Parameters

None

Default

By default evpn mpls multihoming is disabled.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#configure terminal
#(config)# evpn mpls multihoming enable
#(config)# exit
```

evpn mpls vtep-ip-global

Use this command to set the source IP address is used to establish BGP peering with neighbor MPLS Nodes and to establish the EVPN MPLS tunnels.

Use the no version of this command to delete the source IP address.

Command Syntax

```
evpn mpls vtep-ip-global A.B.C.D  
no evpn mpls vtep-ip-global A.B.C.D
```

Parameters

A.B.C.D

IPv4 address type.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
(config)#evpn mpls vtep-ip-global 10.10.11.1  
(config)#no evpn mpls vtep-ip-global 10.10.11.1
```

evpn multi-homed

Use this command to configure interfaces as multi-homed and configure esi-value in case of physical and static lag and system-mac in case of Dynamic lag.

The command `evpn multi-homed` allows users to configure single-active and port-active load-balancing Ethernet Segment Identifier (ESI) on a link with a multihomed Customer Edge (CE).

Use the `no` parameter of this command to unconfigure multi-homed on the interface.

Command Syntax

```
evpn multi-homed esi XX:XX:XX:XX:XX:XX:XX:XX (load-balancing (single-active) |)
evpn multi-homed system-mac (XX-XX-XX-XX-XX-XX|XX:XX:XX:XX:XX:XX|XXXX.XXXX.XXXX) (load-balancing
(port-active) |)
no evpn multi-homed (esi | system-mac)
```

Parameters

XX:XX:XX:XX:XX:XX:XX:XX

ESI value in HH:HH:HH:HH:HH:HH:HH:HH - 9 octet format

XX-XX-XX-XX-XX-XX

Host MAC address (Option 1)

XX:XX:XX:XX:XX:XX

Host MAC address (Option 2)

XXXX.XXXX.XXXX

Host MAC address (Option 3)

port-active

To support ESI configuration for port-active redundancy mode.

single-active

To support ESI configuration for single-active redundancy mode.

Default

Default value is 0.

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 4.0 and underwent modifications in the OcNOS version 6.4.1 and OcNOS version 6.4.2.

Example

```
#configure terminal
(config)#interface xe1
(config)#evpn multi-homed esi 00:11:22:33:44:55:66:77:88 load-balancing single-active
(config)#exit

configure terminal
```

```
(config)#interface po1
(config)#evpn multi-homed system-mac 0000.0000.1111 load-balancing port-active
(config)#exit
```


hardware-profile filter evpn-mpls-cw

This hardware filter needs to be enabled, to allow configuring EVPN-MPLS with control-word functionality per EVI. Filter cannot be disabled if "EVPN-MPLS" is globally enabled.

Before enabling EVPN MPLS per-evi control-word ([evpn mpls id \(page 1936\)](#) command), give this command.

Before disabling the hardware-profile, disable EVPN MPLS globally (no [evpn mpls enable \(page 1554\)](#)).



Note: You need to save the configuration and do a reboot after giving this command.

Use the disable form of this command to disable the configured hardware-profile.

Command Syntax

```
hardware-profile filter evpn-mpls-cw enable
hardware-profile filter evpn-mpls-cw disable
```

Parameters

None

Default

By default, EVPN MPLS control-word hardware-profile is disabled.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.1.0.

Example

```
#configure terminal
(config)# hardware-profile filter evpn-mpls-cw enable
(config)# hardware-profile filter evpn-mpls-cw disable
```

hardware-profile filter evpn-mpls-mh

Use this command to enable the hardware-profile for EVPN MPLS multihoming to successfully activate multihoming in the hardware.

Before enabling EVPN MPLS multihoming ([evpn mpls multihoming enable \(page 1561\)](#) command), give this command.

Before disabling the hardware-profile, disable EVPN MPLS multihoming.



Note: You need to save the configuration and do a reboot after giving this command.

Use the `disable` form of this command to disable the configured hardware-profile.

Command Syntax

```
hardware-profile filter evpn-mpls-mh enable
hardware-profile filter evpn-mpls-mh disable
```

Parameters

None

Default

By default, EVPN MPLS multihoming hardware-profile is disabled.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
#configure terminal
(config)# hardware-profile filter evpn-mpls-mh enable
(config)# hardware-profile filter evpn-mpls-mh disable
```

hardware-profile statistics evpn-mpls

Use this command to enable or disable filter statistics in hardware for evpn mpls network counters.



Note: You need to save the configuration and do a reboot after giving this command.

Command Syntax

```
hardware-profile statistics evpn-mpls enable
hardware-profile statistics evpn-mpls disable
```

Parameters

None

Default

By default, evpn mpls statistics profile is disabled.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
#configure terminal
(config)# hardware-profile statistics evpn-mpls enable
(config)# hardware-profile statistics evpn-mpls disable
```

host-reachability-protocol

Use this command to set the host reachable protocol to Ethernet-VPN over BGP. This defines BGP as the mechanism for host reachability advertisement to discover EVPN peers and to learn remote host details.

Use no form of this command to remove Ethernet-VPN as the host reachable protocol.

Command Syntax

```
host-reachability-protocol evpn-bgp WORD
no host-reachability-protocol evpn-bgp
```

Parameters

WORD

MAC Routing/Forwarding instance name. Maximum limit 32 characters

Default

None

Command Mode

EVPN MPLS Config mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
(config)#evpn mpls id 1
(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf1
```

evpn-route-map

Use this command to configure route-map on an evpn instance.

Use no form of this command to unconfigure route-map from an evpn instance.



Note: This command imports the extended-community color value configured on the route-map.

Command Syntax

```
evpn-route-map WORD
no evpn-route-map
```

Parameters

WORD

Name of the route-map.

Default

None

Command Mode

EVPN MPLS Config mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Examples

```
OcNOS(config)#evpn mpls id 1
OcNOS(config-evpn-mpls)#evpn-route-map map1
OcNOS(config-evpn-mpls)#

OcNOS(config)#evpn mpls id 1
OcNOS(config-evpn-mpls)#no evpn-route-map
OcNOS(config-evpn-mpls)#s
```

ip address

Use this command to set anycast flag for primary and secondary subnets under IRB interface.

With this anycast gateway can be supported for multiple subnets.

Command Syntax

```
ip address [ <A.B.C.D/M> | anycast]
ip address [ <A.B.C.D> | <A.B.C.D> | anycast]
ip address [ <A.B.C.D/M> | secondary | anycast]
ip address [ <A.B.C.D> | <A.B.C.D> | secondary | anycast]
```

Default

The default value is router mac.

Parameters

anycast

Anycast flag

secondary

Used for secondary address option

Command Mode

IRB_IF Mode

Applicability

The anycast flag was introduced in OcNOS version 6.3.0.

Example

```
(config)#interface irb2
(config-irb-if)#ip address 40.1.1.1/24 anycast
(config-irb-if)#ip address 41.1.1.1/24 secondary anycast
(config-irb-if)#
(config)#interface irb1
(config-irb-if)#ip address 42.1.1.1 255.255.255.0 anycast
(config-irb-if)#ip address 43.1.1.1 255.255.255.0 secondary anycast
(config-irb-if)#
```

ipv6 address

Use this command to set anycast flag for any configured subnets under IRB interface.

With this anycast gateway can be supported for multiple subnets.

Command Syntax

```
ipv6 address [ < X:X::X:X/M > | anycast]
```

Default

The default value is router mac.

Parameters

anycast

Anycast flag

Command Mode

IRB_IF Mode

Applicability

The anycast flag was introduced in OcNOS version 6.3.0.

Example

```
(config)#interface irb1
(config-irb-if)# ipv6 address 1100::1/64 anycast
```

link-loss forwarding

Use this command to configure the link-loss forwarding on an access-interface.

Use the `no` parameter to un-configure the load interval globally.

Command Syntax

```
llf-enable  
no llf-enable
```

Parameters

None

Default

None

Command Mode

Access-if mode

Applicability

This command was introduced in OcNOS version 6.6.0.

Example

```
#configure terminal  
(config)#int xe22.2 sw  
(config-if)#encapsulation dot1q 10  
(config-if)#access-if-evpn  
(config-acc-if-evpn)#llf-enable
```


mac-holdtime

Use this command to set the MAC hold time for a MAC/IP or MAC.

The feature holds the MAC in hardware until BGP has withdrawn from the neighbors. This helps to reduce the flooding to other access ports.

This setting applies when the L2 Subifp is shut down, the physical port on which the access port is down, or the access port is removed from the VNID using the no form of the map vnid command.

When the MAC hold time is configured as -1, then the MAC is not removed from the hardware and is also not withdrawn from EVPN BGP.

Use the no form of this command to remove the MAC hold time for the MAC/IP or MAC.



Note: When a MAC is moved to discard state, traffic to and from this MAC is discarded. This is applicable only on statically configured MAC/MAC-IPs.

Command Syntax

```
mac-holdtime <-1-300>  
no mac-holdtime
```

Parameters

<-1-300>

MAC hold time in seconds. Specify -1 to "never expire".

Default

The default holdtime for mac is 3 seconds.

Command Mode

EVPN MPLS mode and ACC_IF mode



Note: When configured in both modes, then the ACC_IF mode value takes preference for that access port.

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
#configure terminal  
(config)#evpn mpls id 3  
(config-evpn-mpls)#mac-holdtime -1  
(config-evpn-mpls)#exit
```

mac-vrf

Use this command to create a MAC VRF to use in EVPN routes.

Use the `no` form of this command to delete the MAC VRF.

Command Syntax

```
mac vrf WORD
no mac vrf WORD
```

Parameters

WORD MAC

Routing/Forwarding instance name. Maximum limit 32 characters

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
#configure terminal
(config)#mac vrf vrf1
(config)#no mac vrf vrf1
```

map vpn-id

Use this command to map a sub-interface to a tenant.

Use the no form of this command to remove the tenant.

Command Syntax

```
map vpn-id <1-16777215> (etree-leaf|)
no map vpn-id <1-16777215>
```

Parameters

vpn-id <1-16777215>

Specifies the EVPN Instance (EVI) ID (VNID) to which the interface belongs.

etree-leaf

(Optional) Enables configuring the access interface (AC) or sub-interface as a Leaf in an EVPN E-Tree deployment. This allows the PE to enforce E-Tree traffic rules, including Leaf-to-Leaf traffic blocking and Root-to-Leaf communication.

Command Mode

Access interface EVPN mode

Applicability

Introduced in OcNOS version 3.0. Added `etree-leaf` parameter in OcNOS version 7.0.0.

Examples

1. Configure a regular EVPN interface (no E-Tree Leaf):

```
#configure terminal
(config)#interface xel.1 switchport
(config-if)#access-if-evpn
(config-access-if)#map vpn-id 1
(config-access-if)#end
```

2. Configure an interface as an E-Tree Leaf AC:

```
#configure terminal
(config)#evpn etree enable scenario-2
(config)#evpn mpls id 203 etree-leaf
(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf103
(config-evpn-mpls)#exit

(config)#interface xel.1 switchport
(config-if)#access-if-evpn
(config-access-if)#map vpn-id 3 etree-leaf
(config-access-if)#end
```

nd-cache disable

Use this command to disable ND cache for MAC/IPv6.

When the ND cache is disabled on an Access Circuits, It does not reply to any ND arriving on this port from the cache. It withdraws all MAC/IPs configured/learned on this Access Circuit and removes the MAC/IP entry for this AC from the local ND cache.

It also makes sure that on withdrawing the MAC/IP route, the MAC does not become unknown. If all routes for this MAC are being withdrawn because of this command, then it advertises a MAC-only route. This is done so that the MAC does not become unknown and only the cache functionality becomes disabled.

Use the `no` form of this command to enable ND cache for MAC/IPv6.



Notes:

- On enabling the cache, an IP will be in conflict, then the cache enable will fail. The conflict has to be manually removed and then the cache enabled.
- Not applicable for the AC port which is mapped with ELINE/Xconnect Service.

Command Syntax

```
nd-cache disable
no nd-cache disable
```

Parameters

None

Default

By default, the nd-cache option is enabled.

Command Mode

Access if evpn mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
#configure terminal
(config)#interface xel.1 switchport
(config-if)#access-if-evpn
(config-access-if)#nd-cache disable
(config-access-if)#end
```

rewrite egr-pcp-disable

Use this command to disable egress QoS remarking on a Layer 2 Attachment Circuit (AC)

Use the `no` form of this command to enable traffic remarking.



Note: In L2vpn scenario, it preserves customer VLAN priority bits by disabling egress QoS remarking on the Layer 2 Attachment Circuit (AC)

Command Syntax

```
rewrite egr-pcp-disable  
no rewrite egr-pcp-disable
```

Parameters

None

Default

By default, the traffic remarking is enabled.

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 6.6.1.

Example

```
#configure terminal  
(config)#interface xe8.1 switchport  
(config-if)#encapsulation dot1q 2346  
(config-if)#rewrite egr-pcp-disable  
(config-if)#mtu 9100  
(config-if)#access-if-evpn  
(config-access-if)#map vpn-id 729  
(config-access-if)#end
```

service-carving

Use this command to provide the flexibility to select the Designated Forwarder (DF) election algorithm based on preference based or modulo-based DF election.

Use no form of this command to disable service-carving.

Command Syntax

```
service-carving (preference-based|auto)
no service-carving
```

Parameters

preference-based

Select the DF election algorithm based on preference based.

auto

Select the DF election algorithm based on modulo based.

Default

None

Command Mode

EVPN ES Mode

Applicability

This command was introduced in the OcNOS version 6.4.1.

Example

The following examples demonstrate the configuration of the `service-carving` command in both `single-active` or `port-active` mode for the EVPN multi-homed system, with one utilizing `auto` service carving and the other using `preference-based` service carving.

```
OcNOS#configure terminal
OcNOS(config)#interface sal
OcNOS(config-if)#evpn multi-homed esi 11:22:33:44:55:66:77:88:99 load-balancing single-active
OcNOS(config-if-es)#service-carving auto
OcNOS(config-if-es)#end

OcNOS#configure terminal
OcNOS(config)#interface pol
OcNOS(config-if)#evpn multi-homed system-mac 0000.0000.0011 load-balancing port-active
OcNOS(config-if-es)#service-carving auto
OcNOS(config-if-es)#end

OcNOS#configure terminal
OcNOS(config)#interface pol
OcNOS(config-if)#evpn multi-homed system-mac 0000.0000.0011 load-balancing port-active
OcNOS(config-if-es)#service-carving preference-based
OcNOS(config-if-es)#end
```

The following example is used to disable the `service-carving` for the EVPN multi-homed system.

```
OcNOS (config-if-es) #no service-carving  
OcNOS (config-if-es) #end
```

service-carving ac-driven

Use this command to enable the AC-influenced method for any selected Designated Forwarder (DF) algorithm.

Enabling the `ac-driven` method allows the Designated Forwarder (DF) algorithm to be influenced by the Attachment Circuits (ACs) associated with a specific Ethernet Segment (ES). This means that the DF selection is based on the ACs characteristics and conditions, such as whether an AC is operational UP, mapped, or unmapped on the ESI.

Use `no` form of this command to disable the AC-influenced method for any selected Designated Forwarder (DF) algorithm.

Command Syntax

```
service-carving ac-driven
no service-carving ac-driven
```

Parameters

None

Default

`ac-driven` is disabled.

Command Mode

EVPN Ethernet Segment (ES) Mode

Applicability

This command was introduced in the OcNOS version 6.4.2.

Example

The provided examples showcase the configuration of the `service-carving ac-driven` command in EVPN Ethernet Segment (ES) mode. The first two examples demonstrate enabling this feature with different DF election methods, and the final example illustrates the command to disable `service-carving ac-driven`.

```
OcNOS#configure terminal
OcNOS(config)#interface sal
OcNOS(config-if)#evpn multi-homed esi 11:22:33:44:55:66:77:88:99 load-balancing single-active
OcNOS(config-if-es)#service-carving preference-based
OcNOS(config-if-es)#service-carving ac-driven
OcNOS(config-if-es)#end
```

```
OcNOS#configure terminal
OcNOS(config)#interface sal
OcNOS(config-if)#evpn multi-homed esi 11:22:33:44:55:66:77:88:99 load-balancing single-active
OcNOS(config-if-es)#service-carving auto
OcNOS(config-if-es)#service-carving ac-driven
```

```
OcNOS(config-if-es)#no service-carving ac-driven
OcNOS(config-if-es)#end
```


service-carving weight

Use this command to specify a preference value when the preference-based Designated Forwarder (DF) election algorithm is selected. This preference value determines the priority of the local PE device to become the DF for a particular Ethernet segment.

Use no form of this command to replace the preference weight value and choose the default preference value.

Command Syntax

```
service-carving weight <1-65535>  
no service-carving weight
```

Parameters

weight <1-65535>

Specifies the preference weight value. A lower weight value indicates a higher priority for becoming the DF.

Default

The service-carving weight command is set to 32767 by default.

Command Mode

EVPN Ethernet Segment (ES) Mode

Applicability

This command was introduced in the OcNOS version 6.4.1.

Example

The `service-carving weight` command is used to configure the preference weight value for service-carving in both port-active and single-active modes.

```
OcNOS#configure terminal  
OcNOS(config)#interface pol  
OcNOS(config-if)#evpn multi-homed system-mac 0000.0000.0011 load-balancing port-active  
OcNOS(config-if-es)#service-carving preference-based  
OcNOS(config-if-es)#service-carving weight 100  
OcNOS(config-if-es)#end  
  
OcNOS#configure terminal  
OcNOS(config)#interface sal  
OcNOS(config-if)#evpn multi-homed esi 11:22:33:44:55:66:77:88:99 load-balancing single-active  
OcNOS(config-if-es)#service-carving preference-based  
OcNOS(config-if-es)#service-carving weight 100
```

To disable the configured weight, use the `no service-carving weight` command.

```
OcNOS(config-if-es)#no service-carving weight  
OcNOS(config-if-es)#end
```

show bgp l2vpn evpn

Use this command to display details about Layer 2 Virtual Private Network (L2VPN) Ethernet Virtual Private Network (EVPN) routes.



Notes:

- A BGP EVPN route update received for an unreachable IP address is also listed by this command and as a best route. This is because the next hop tracking feature is not supported for the EVPN address family. However, the tunnel to this IP address is shown in unresolved state by the show nvo vxlan tunnel output.
- An E-tag (Ethernet tag) can have the value of zero/VID/VNID based on the use case. An E-tag can go up to 32 bits and no restrictions are noted in the RFC. Since an E-tag can have different values, it should not be compared with the label/VNID.

Command Syntax

```
show bgp l2vpn evpn ((vrf WORD)|(rd WORD)| time|)
show bgp l2vpn evpn mac-ip ((vrf WORD)|(rd WORD)|)
show bgp l2vpn evpn mcast
show bgp l2vpn evpn multihoming es-route <(rd WORD)|(vrf WORD)>
show bgp l2vpn evpn multihoming ethernet-ad-per-evi <(rd WORD)|(vrf WORD)>
show bgp l2vpn evpn multihoming ethernet-ad-per-es <(rd WORD)|(vrf WORD)>
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
MH-PE2#sh bgp l2vpn evpn
BGP table version is 22, local router ID is 102.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

      Network          Next Hop          Metric      LocPrf      Weight      Path  Peer      Encap
RD[101.1.1.1:1]
*>i  [1]:[00:00:00:00:00:15:15:00:00:00]:[6001]:[18]
```

```

101.1.1.1          0          100          0          i 101.1.1.1      MPLS
*>i  [3]:[6001]:[32,101.1.1.1]
101.1.1.1          0          100          0          i 101.1.1.1      MPLS

RD[101.1.1.1:2]
*>i  [1]:[00:00:00:00:00:15:15:00:00:00]:[10]:[19]
101.1.1.1          0          100          0          i 101.1.1.1      MPLS

RD[101.1.1.1:22]
*>i  [1]:[00:00:00:00:00:15:15:00:00:00]:[4294967295]:[0]
101.1.1.1          0          100          0          i 101.1.1.1      MPLS
*>i  [4]:[00:00:00:00:00:15:15:00:00:00]:[32,101.1.1.1]
101.1.1.1          0          100          0          i 101.1.1.1      MPLS

RD[102.1.1.1:1] VRF[l2vrf1]:
*>  [1]:[00:00:00:00:00:15:15:00:00:00]:[6001]:[19]
102.1.1.1          0          100          32768      i -----      MPLS
* i  101.1.1.1          0          100          0          i 101.1.1.1      MPLS
* i  [1]:[00:00:00:00:00:15:15:00:00:00]:[4294967295]:[0]
101.1.1.1          0          100          0          i 101.1.1.1      MPLS
*>  [2]:[00:00:00:00:00:15:15:00:00:00]:[6001]:[48,0000:0000:aa11]:[0]:[19]
102.1.1.1          0          100          32768      i -----      MPLS
*>  [2]:[00:00:00:00:00:15:15:00:00:00]:[6001]:[48,0000:0000:aa12]:[32,12.12.12.10]:[19]
102.1.1.1          0          100          32768      i -----      MPLS
*>  [2]:[00:00:00:00:00:15:15:00:00:00]:[6001]:[48,0000:0000:aa13]:[128,1201::10]:[19]
102.1.1.1          0          100          32768      i -----      MPLS
* i  [3]:[6001]:[32,101.1.1.1]
101.1.1.1          0          100          0          i 101.1.1.1      MPLS
*>  [3]:[6001]:[32,102.1.1.1]
102.1.1.1          0          100          32768      i -----      MPLS
* i  [3]:[6001]:[32,105.1.1.1]
105.1.1.1          0          100          0          i 105.1.1.1      MPLS

RD[102.1.1.1:2] VRF[l2vrf2]:
* i  [1]:[0]:[11]:[641]
105.1.1.1          0          100          0          i 105.1.1.1      MPLS
*>  [1]:[00:00:00:00:00:15:15:00:00:00]:[10]:[18]
102.1.1.1          0          100          32768      i -----      MPLS
* i  101.1.1.1          0          100          0          i 101.1.1.1      MPLS
* i  [1]:[00:00:00:00:00:15:15:00:00:00]:[4294967295]:[0]
101.1.1.1          0          100          0          i 101.1.1.1      MPLS

RD[102.1.1.1:22] VRF[evpn-gvrf-1]:
*>  [1]:[00:00:00:00:00:15:15:00:00:00]:[4294967295]:[0]
102.1.1.1          0          100          32768      i -----      MPLS
* i  [4]:[00:00:00:00:00:15:15:00:00:00]:[32,101.1.1.1]
101.1.1.1          0          100          0          i 101.1.1.1      MPLS
*>  [4]:[00:00:00:00:00:15:15:00:00:00]:[32,102.1.1.1]
102.1.1.1          0          100          32768      i -----      MPLS

RD[105.1.1.1:1]
*>i  [3]:[6001]:[32,105.1.1.1]
105.1.1.1          0          100          0          i 105.1.1.1      MPLS

RD[105.1.1.1:2]
*>i  [1]:[0]:[11]:[641]
105.1.1.1          0          100          0          i 105.1.1.1      MPLS

Total number of prefixes 21
MH-PE2#
MH-PE2#

```

show evpn load-balance

Use this command to display ESI information in Single-Active or Port-Active mode. It assists the user in identifying whether the local device is in Active or Standby mode and understanding which election algorithm is used for ESI. In remote devices, information will only be displayed for ELINE services.

This command is used to debug and understand if the election process is occurring correctly. The commands on multi-homed devices are symmetric, and they provide insights into the election algorithm used for the DF election.

Command Syntax

```
show evpn load-balance (port-active | single-active | all)
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced before OcNOS version 6.4.1 and underwent modifications in the OcNOS version 6.4.2

Example

The following example displays ESI information using the `show evpn load-balance port-active` and `show evpn load-balance single-active` commands.

```
OcNOS#show evpn load-balance port-active
ESI          AC-IF/PE  PE-IP-ADDRESS  Redundancy  Service-
carving      weight  Revertive      AC-DF      Status
=====
00:00:00:00:00:12:12:00:00:00  REMOTE  101.1.1.1      port-
active      auto    0              NO          NA  ACTIVE
00:00:00:00:00:12:12:00:00:00  LOCAL   102.1.1.1      port-
active      auto    0              NO          NA  STANDBY

#show evpn load-balance single-active
ESI          AC-IF/PE  PE-IP-ADDRESS  Redundancy  Service-
carving      weight  Revertive      AC-DF      Status
=====
00:11:22:33:00:00:00:55:66:77  sa1.1    1.2.3.4        single-active  preference-
based      200     NO             NO          ACTIVE
00:11:22:33:00:00:00:55:66:77  ----    3.4.5.6        single-active  preference-
based      100     NO             NO          ----
```

The following table explains the output fields.

Table 84. show evpn load-balance port-active

| Field | Description |
|-----------------|--|
| ESI | Ethernet Segment Identifier, a unique identifier for an Ethernet Segment. |
| AC-IF/PE | This field indicates whether the ESI is associated with a local or remote Attachment Circuit (AC) or Provider Edge (PE). |
| PE-IP-ADDRESS | The IP address of the PE associated with the ESI. |
| Redundancy | Indicates whether the ESI is configured for port-active or single-active redundancy. |
| Service-carving | The service carving mode associated with the ESI. |
| weight | Weight assigned to the ESI. |
| Revertive | Indicates whether the ESI is configured for revertive mode. |
| AC-DF | Attachment Circuit Designated Forwarder status. |
| Status | The status of the ESI, whether it's active or standby. |

show evpn mpls

Use this command to display the EVPN Information.

Command Syntax

```
show evpn mpls ([id <1-16777215>])
```

Parameters

<1-16777215>

EVPN-MPLS tenant identifier

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
MH-PE2#sh evpn mpls
EVPN-MPLS Information
=====
Codes: NW - Network Port
       AC - Access Port
       (u) - Untagged

VPN-ID   EVI-Name      EVI-Type Type Interface ESI              VLAN DF-Status Src-
Addr          Dst-Addr

6001     ----         L2      NW    ----      ----          ----  ----
        102.1.1.1      101.1.1.1
6001     ----         L2      NW    ----      ----          ----  ----
        102.1.1.1      105.1.1.1
6001     ----         --      AC    xe1.301    --- Single Homed Port ---  ----  ----
        ----
6001     ----         --      AC    po10.301   00:00:00:00:00:15:15:00:00:00  ----  DF      ----
        ----
6002     ----         L2      NW    ----      ----          ----  ----
        102.1.1.1      101.1.1.1
6002     ----         L2      NW    ----      ----          ----  ----
        102.1.1.1      105.1.1.1
6002     ----         --      AC    po10.302   00:00:00:00:00:15:15:00:00:00  ----  NON-DF  ----
        ----

Total number of entries are 7
MH-PE2#
```

show evpn mpls arp-cache

Use this command to display the ARP cache information.

Command syntax

```
show evpn mpls arp-cache (|evid <1-16777215>) summary
```

Parameters

evid

EVPN-MPLS tenant identifier

<1-16777215>

Range supported for EVID

summary

Ethernet Virtual Private Network ID

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
PE1#show evpn mpls arp-cache
MPLS-EVPN ARP-CACHE Information
=====
EVPN-ID   Ip-Addr      Mac-Addr      Type      Age-Out   Retries-Left
-----
2001      122.122.122.1  0000.0000.1111 Dynamic Local   ----
2001      122.122.122.2  0000.0000.2222 Dynamic Remote  ----
Total number of entries are 2
PE1#
```

show evpn mpls counters

This command to display the ingress/egress evpn-mpls network port statistic counters for both BUM and known unicast data traffic.



Note: To see the statistics, you must enable the command `hardware-profile statistics evpn-mpls enable` and reboot the board for the command to take effect.

Command syntax

```
show evpn mpls counters ((network (ingress|egress dst <A.B.C.D>))) | vpn-id <1-16777215> (network (ingress|egress dst <A.B.C.D>)))
```

Parameters

<1-16777215>

VPN identifier

A.B.C.D

IPv4 address

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
DUT2#show evpn mpls counters network egress
+-----+-----+-----+
| VPN-ID | DESTINATION | BUM | Unicast |
|        | PEER        | TX (pkts) | TX (pkts) |
+-----+-----+-----+
| 2000    | 2.2.2.2     | 140136 | 68900   |
| 1999    | 2.2.2.2     | 140138 | 68899   |
| 1998    | 2.2.2.2     | 140136 | 68899   |
| 1997    | 2.2.2.2     | 140137 | 68897   |
DUT2#
DUT2#show evpn mpls counters network ingress
+-----+-----+-----+
| VPN-ID | BUM | Unicast |
|        | RX (pkts) | RX (pkts) |
+-----+-----+-----+
| 30     | 0   | 800122  |
| 10     | 0   | 727383  |
| 4001   | 5   | 727362  |
| 3003   | 0   | 0       |
| 3002   | 727368 | 0       |
| 3001   | 3   | 727364  |
DUT2#
```


show evpn mpls label alias

Use this command to display the alias label details of the MH Nodes.

Command syntax

```
show evpn mpls label alias
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
Remote-PE3#sh evpn mpls label alias
S - Self
R - Remote
ESI                               PE-IP-ADDRESS      TENANT      ALIAS-LABEL
=====
00:00:00:00:00:15:15:00:00:00    101.1.1.1 (R)      6001        18
00:00:00:00:00:15:15:00:00:00    101.1.1.1 (R)      6002        20
00:00:00:00:00:15:15:00:00:00    101.1.1.1 (R)      11          19
00:00:00:00:00:15:15:00:00:00    102.1.1.1 (R)      6001        19
00:00:00:00:00:15:15:00:00:00    102.1.1.1 (R)      6002        20
00:00:00:00:00:15:15:00:00:00    102.1.1.1 (R)      11          18
Remote-PE3#
Remote-PE3#
```

show evpn mpls label esi

Use this command to display the esi label details of MH Nodes.

Command syntax

```
show evpn mpls label esi
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.0.

Examples

```
MH-PE1#sh evpn mpls label esi
S - Self
R - Remote
ESI                               PE-IP-ADDRESS          ESI-LABEL
=====
00:00:00:00:00:13:13:00:00:00    101.1.1.1 (S)           38
00:00:00:00:00:13:13:00:00:00    102.1.1.1 (R)           43
00:00:00:00:00:15:15:00:00:00    101.1.1.1 (S)           37
00:00:00:00:00:15:15:00:00:00    102.1.1.1 (R)           42
```

show evpn mpls mac-table

Use this command to display the host MAC address table.

Command Syntax

```
show evpn mpls mac-table ([evnid <1-16777215>|]) (summary | hardware |)
```

Parameters

<1-16777215>

EVPN-MPLS tenant identifier

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0 and enhanced in OcNOS version 7.0.0 to show the output in ascending order of VNID.

Example

```
DUT1#show evpn mpls mac-table
```

```
=====
EVPN MPLS MAC Entries
=====
```

| VNID Ip/ESI | Interface | VlanId | In-VlanId | Mac-Addr Type | VTEP- Status | MAC move | AccessPortDesc | LeafFlag |
|------------------------|-----------|--------|-----------|---------------------------------|-----------------|----------|----------------|----------|
| 501 25.0.1.1 -- | xe7.500 | ---- | ---- | 9819.2c47.2809 Dynamic Local | ----- | 300 | ----- | -- |
| 701 25.0.1.1 -- | xe48.700 | ---- | ---- | e49d.73d2.5b1d Dynamic Local | ----- | 300 | ----- | -- |
| 3000 25.0.1.1 -- | xe7.3000 | ---- | ---- | 9819.2c47.2809 Dynamic Local | ----- | 300 | ----- | -- |
| 3500 25.0.1.1 -- | xe48.3500 | ---- | ---- | e49d.73d2.5b1d Static Local | ----- | 300 | ----- | -- |

Total number of entries are : 4

```
DUT1#
```

```
DUT1#show evpn mpls mac-table hardware
```

```
=====
EVPN MPLS MAC Entries
=====
```

| VNID | Interface | VlanId | In-VlanId | Mac-Addr | VTEP- |
|------|-----------|--------|-----------|----------|-------|
|------|-----------|--------|-----------|----------|-------|

| Ip/ESI | | | | Type | Status | Time-out | AccessPortDesc |
|----------|-----------|------|------|----------------|--------|----------|----------------|
| | | | | | | | |
| 501 | xe7.500 | ---- | ---- | 9819.2c47.2809 | | | |
| 25.0.1.1 | | | | Dynamic Local | ----- | 300 | ----- |
| 701 | xe48.700 | ---- | ---- | e49d.73d2.5b1d | | | |
| 25.0.1.1 | | | | Dynamic Local | ----- | 300 | ----- |
| 3000 | xe7.3000 | ---- | ---- | 9819.2c47.2809 | | | |
| 25.0.1.1 | | | | Dynamic Local | ----- | 300 | ----- |
| 3500 | xe48.3500 | ---- | ---- | e49d.73d2.5b1d | | | |
| 25.0.1.1 | | | | Static Local | ----- | 300 | ----- |

```
Total number of entries are 4
DUT1#
DUT1#show evpn mpls mac-table summary
=====
Evpn mpls MAC Summary
=====

Total number of entries are : 4

DUT1#
```

show evpn mpls nd-cache

Use this command to display the ND cache information.

Command Syntax

```
show evpn mpls nd-cache ([|evid <1-16777215>) summary
```

Parameters

evid

EVID

<1-16777215>

Range supported for EVID

summary

Ethernet Virtual Private Network ID

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
PE1#
PE1#show evpn mpls nd-cache
MPLS-EVPN ND-CACHE Information
=====
EVPN-ID   Ip-Addr                               Mac-Addr                               Type           Age-Out   Retries-Left
-----
1001      1201::1                             0000.0000.cc10 Dynamic Local   ----
1001      1201::2                             0000.0000.dd10 Dynamic Remote  ----
Total number of entries are 2
PE1#
PE1#
```

show evpn mpls route-count

Use this command to display the evpn active route (MAC-IP,MAC-IPv6 and MAC-only) count information.

Command Syntax

```
show evpn mpls route-count (|evid <1-16777215>)
```

Parameters

evid

EVPN-MPLS tenant identifier

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 5.0 and enhanced in OcNOS version 7.0.0 to show the output in ascending order of VNID.

Examples

```
DUT1#show evpn mpls route-count
EVPN-MPLS Active route count information
=====
Max supported route count   : 131072

-----
VNID      Total      MACONLY  MACIPv4  MACIPv6
-----
501        1          1         0         0
701        1          1         0         0
3000       1          1         0         0
3500       1          1         0         0

Active route count: 4
Total number of entries are 4
DUT1#
```

The following table explains the output fields.

Table 85. show evpn mpls route-count fields

| Field | Description |
|--------------------|---|
| Max route count | Maximum number of route count in evpn mpls. |
| Active route count | Number of active route count in the interface. |
| VNID | VNID is used to identify Layer 2 segments and to maintain Layer 2 isolation between the segments. |
| Total | Total number of entries for the interface. |
| MACONLY | The MAC-only route for both the local and the remote entries in the EVPN MPLS |

Table 85. show evpn mpls route-count fields (continued)

| Field | Description |
|---------|--|
| | instance route table. |
| MACIPv4 | IPv4 media access control (MAC) address for a default virtual gateway. |
| MACIPv6 | IPv6 media access control (MAC) address for a default virtual gateway. |

show evpn mpls static host state

Use this command to display the state of the host which is configured statically.

Command Syntax

```
show evpn mpls static host state
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
MH-PE2#sh evpn mpls static host status
```

```
MPLS Static Host Information
```

```
=====
```

```
Codes: NW - Network Port
```

```
AC - Access Port
```

```
(u) - Untagged
```

| VNID Addr | Ifname Status | Outer-Vlan | Inner-vlan | Ip-Addr | Mac- |
|--------------|------------------|------------|------------|-------------|----------------|
| 6001 | po10.301 | --- | --- | 0.0.0.0 | 0000.0000.aa11 |
| Active | | | | | |
| 6001 | po10.301 | --- | --- | 12.12.12.10 | 0000.0000.aa12 |
| Active | | | | | |
| 6001 | po10.301 | --- | --- | 1201::10 | 0000.0000.aa13 |
| Active | | | | | |

```
Total number of entries are 3
```

```
MH-PE2#
```


show evpn mpls status

Use this command to display EVPN MPLS status in hardware.

Command Syntax

```
show evpn mpls status
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
PE-1#show evpn mpls status  
EVPN-MPLS is ACTIVE in Hardware
```

show evpn mpls tunnel

Use this command to view the active EVPN MPLS tunnels for ELAN.

Command Syntax

```
show evpn mpls tunnel [| summary]
```

Parameters

summary

Summary information

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Example

```
PE1#sh evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status          Up/Down          Update           evpn-id
=====
101.1.1.1        105.1.1.1       Installed       00:43:09         00:43:09         2001
101.1.1.1        105.1.1.1       Installed       00:43:09         00:43:09         1001

Total number of entries are 2
PE1#
```

show evpn mpls tunnel label

Use this command to display the label details of the EVPN tunnels.

Command Syntax

```
show evpn mpls tunnel label
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0 and RSVP multipath was introduced in OcNOS version 5.0.

Example

```
PE1#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels
(*) in Policy - tunnel-policy inherited from mac-vrf
=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====
VP-Multipath      Underlay
Destination      Status      VPN-ID      Policy
Name      NHLFE-ix  NW-Intf      NW-Label
=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====
3.3.3.3      --      Installed    601      --
              --      xe10          24320
3.3.3.3      --      Installed    801      --
              --      xe10          24320
4.4.4.4      --      Installed    1601     --
              642      20          640      17      MP1      12      NA      NA
4.4.4.4      --      Installed    1801     --
              643      21          641      18      MP1      12      NA      NA

Total number of entries are 4
PE1#
```

show evpn mpls xconnect

Use this command to display the VPWS xconnect details of the MTU, AC-NW connections, and network status.

Command Syntax

```
show evpn mpls xconnect
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
Remote-PE3#sh evpn mpls xconnect
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                               Remote      Connection-Details
=====
=====
VPN-ID      EVI-Name      MTU  VPN-ID      Source      Destination      PE-
IP          MTU   Type   NW-Status
=====
11          ----          1500  10          xe1.10      00:00:00:00:00:15:15:00:00:00
102.1.1.1   1500  AC-NW  NW-SET
                                     101.1.1.1
1500 ---- ----
40          ----          1500  30          xe2.2      --- Single Homed Port ---
102.1.1.1   1500  AC-NW  NW-SET

Total number of entries are 2
Remote-PE3#
```

show evpn mpls xconnect tunnel

Use this command to view the active EVPN MPLS Tunnels for ELINE.

Command Syntax

```
show evpn mpls xconnect tunnel [| summary]
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.0.

Example

```
PE1#
PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====
===
101.1.1.1    105.1.1.1        Installed   01:25:48     01:25:48     50           60
101.1.1.1    105.1.1.1        Installed   00:07:53     00:07:53     30           40
101.1.1.1    105.1.1.1        Installed   00:07:55     00:07:55     10           20

Total number of entries are 3
PE1#
```

show evpn mpls xconnect tunnel label

Use this command to display the label details of EVPN tunnels for ELINE.

Command Syntax

```
show evpn mpls xconnect tunnel label
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.0 and RSVP-Multipath is introduced in OcNOS version 5.0.

Example

```
PE1#
PE1#show evpn mpls xco tunnel label
EVPN-MPLS Network tunnel labels
(*) in Policy - tunnel-policy inherited from mac-vrf
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====
Multipath          Underlay      Local      Remote          Local      Remote      RSVP-
Destination      Status      VPWS-ID    VPWS-ID    Policy      UC-Label    UC-Label    Grp-
Name      NHLFE-ix  NW-Intf    NW-Label
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====
3.3.3.3          Installed    501         1           --          NA          NA          --
16              16          MP2         10          NA          19          16          --
4.4.4.4          Installed    1501        1001        --          19          16          --
xe10            24320
Total number of entries are 2
PE1#
```

show running-config evpn mpls

Use this command to display the current running configuration of EVPN MPLS.

Command Syntax

```
show running-config evpn mpls
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0.

Examples

```
#show running-config evpn mpls
!
evpn mpls enable
!
evpn mpls vtep-ip-global 101.1.1.1
!
evpn mpls multihoming enable
!
evpn esi hold-time 30
!
evpn mpls mac-ageing-time 180
!
evpn mpls arp-nd refresh-timer 60
!
evpn mpls id 10 xconnect target-mpls-id 20
  host-reachability-protocol evpn-bgp l2vrf5
!
vpn-id 1001
  host-reachability-protocol evpn-bgp l2vrf1
  mac-holdtime 300
  evi-name ELANE
!
interface xe1.2 switchport
  access-if-evpn
  map vpn-id 1001
  dynamic-learning disable
  arp-nd flood-suppress
  mac-holdtime 200
  arp-cache disable
  nd-cache disable
  mac 0000.0000.aa31
  mac 0000.0000.aa32 ip 12.12.12.10
  mac 0000.0000.aa33 ipv6 1201::1
!
interface xe1.10 switchport
  encapsulation dot1q 10
  access-if-evpn
```

```
map vpn-id 10
!
```

tunnel-select-policy

Use this command to associate tunnel-policy to an EVPN instance and to a mac-vrf. This will be used while selecting transport for EVPN service. In particular, an EVPN instance could be either an ELINE or ELAN type and may also multi-homing also enabled.

Use the `no` parameter with this command to un-configure tunnel-policy from the EVPN instance.

Please refer [MPLS Service Mapping Configuration \(page 1245\)](#) for Tunnel-policy Name configuration.

Command Syntax

```
tunnel-select-policy NAME
no tunnel-select-policy
```

Parameters

NAME

Name of tunnel-policy

Command Mode

EVPN MPLS Config mode and MAC_VRF MODE

Applicability

This command was introduced in OcNOS version 4.1 but the support for this CLI under mac-vrf mode is introduced in OcNOS-SP version 4.2.

Examples

```
#configure terminal
(config)#mac vrf vrf1
(config-vrf)# tunnel-select-policy policy1
(config-vrf)#end
(config)# evpn mpls id 1001
(config-evpn-mpls)# tunnel-select-policy policy2
(config-evpn-mpls)# end
(config)# evpn mpls id 10 xconnect target-mpls-id 20
(config-evpn-mpls)# tunnel-select-policy policy3
(config-evpn-mpls)# end
```

MPLS Service Mapping Commands

This chapter describes each MPLS Service Mapping command.

| | |
|--|------|
| color | 1607 |
| mode | 1608 |
| mpls 6pe tunnel-select-policy | 1609 |
| show running-config tunnel-policy | 1610 |
| show tunnel-policy | 1611 |
| tunnel-policy | 1614 |
| tunnel-select-policy (For L3VPN service) | 1615 |

color

Use this command to set color value for a tunnel policy. This value will be used while selecting transport for given service.

Use the `no` parameter with this command to un-configure color value for tunnel policy.

Command Syntax

```
color <1-4294967295>  
no color <1-4294967295>
```

Parameters

color

Configure color code

Command Mode

Tunnel policy mode

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
#configure terminal  
(config)#tunnel-policy policy1  
(config-tnl-policy)#color 23
```

mode

Use this command to set tunnel selection mode for a tunnel policy. This will be used while selecting transport for given service. Mode can be dedicated or best-effort.

Use the `no` parameter with this command to set mode to default mode.

Command Syntax

```
mode dedicated|best-effort
no mode
```

Parameters

best-effort

Best effort tunnel selection mode

dedicated

Dedicated tunnel selection mode (Default)

Default

Dedicated

Command Mode

Tunnel policy mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
#configure terminal
(config)#tunnel-policy policy1
(config-tnl-policy)# mode best-effort
```

mpls 6pe tunnel-select-policy

Use this command to set tunnel-policy for a 6PE services. This will be used while selecting transport for 6PE services. Use the `no` parameter with this command to unconfigure tunnel selection policy for 6PE.

Command Syntax

```
mpls 6pe tunnel-select-policy NAME
```

Parameters

NAME

Name of tunnel-policy

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
#configure terminal
(config)# mpls 6pe tunnel-select-policy policy1
```

show running-config tunnel-policy

Use this command to display the current running configuration of all tunnel policies.

Command Syntax

```
show running-config tunnel-policy
```

Parameters

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
#show running-config tunnel-policy
!
tunnel-policy tp1
  color 1
!
```

show tunnel-policy

Use this command to display the tunnel policy details.

Command Syntax

```
show tunnel-policy (NAME (service (((6pe|vpls|vpws|evpn|l3vpn|) details)) | (count)))| service details |)
```

Parameters

NAME

show tunnel policy table by name

service

show service information

6pe

Display 6PE service details

evpn

Display EVPN service details

l3vpn

Display L3VPN service details

vpls

Display VPLS service details

vpws

Display VPWS service details

count

show counts

details

show details

service

show service information

details

show details

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

The following sample shows the summary of all configured tunnel policies.

```
#show tunnel-policy  
tunnel-policy tpl
```

```

color 1
tunnel-mode dedicated
tunnel-type any
Total Services attached:
VRF      : 1
VPWS     : 1
VPLS     : 1
EVPN     :

```

The following sample shows the summary of specific configured tunnel policy by name.

```

#show tunnel-policy tp1
tunnel-policy tp1
color 1
tunnel-mode dedicated
tunnel-type any
Total Services attached:
VRF      : 1
VPWS     : 1
VPLS     : 1
EVPN     : 0

```

The following sample shows the detailed information of all the configured tunnel policies.

```

#show tunnel-policy service details
List of services configured with tunnel-policy tp1:
VPWS
  VPWS-Id: 100, Name: vpws1 Status: UP, Mapped Tunnel: RSVP, ftn_ix 4, color 1
L3VPN
  VRF id 2
    FTN id: 1, Prefix: 20.10.30.0/24, Status: Up, Label: 25664, Mapped Tunnel: RSVP, ftn_ix 4, color
1    FTN id: 2, Prefix: 3001::/64, Status: Up, Label: 25664, Mapped Tunnel: RSVP, ftn_ix 4, color 1
VPLS
  VPLS-Id: 300, Name: vpls1, Status: UP, Mapped Tunnel: RSVP, ftn_ix 4, color 1

```

The following sample shows the detailed information of specific configured tunnel policy by name.

```

# show tunnel-policy tp1 service details
Total services using tunnel-policy tp1:
VPWS
  VPWS-Id: 100, Name: vpws1 Status: UP, Mapped Tunnel: RSVP, ftn_ix 4, color 1
L3VPN
  VRF id 2
    FTN id: 1, Prefix: 20.10.30.0/24, Status: Up, Label: 25664, Mapped Tunnel: RSVP, ftn_ix 4, color
1    FTN id: 2, Prefix: 3001::/64, Status: Up, Label: 25664, Mapped Tunnel: RSVP, ftn_ix 4, color 1
VPLS
  VPLS-Id: 300, Name: vpls1, Status: UP, Mapped Tunnel: RSVP, ftn_ix 4, color 1

```

The following sample shows the services count of specific tunnel policy by name.

```

#show tunnel-policy tp1 service count
Total services using tunnel-policy tp1:
VPWS   : 1
VPLS   : 0
L3VPN  : 2
6PE    : 0
EVPN   : 0

```

The following sample shows the detailed VPLS service information of specific tunnel policy by name.

```

#show tunnel-policy tp1 service vpls details
Total services using tunnel-policy tp1:
VPLS
  VPLS-Id: 300, Name: vpls1, Status: UP, Mapped Tunnel: RSVP, ftn_ix 4, color 1

```

The following sample shows the detailed VPWS service information of specific tunnel policy by name.


```
#show tunnel-policy tpl service vpws details
Total services using tunnel-policy tpl:
VPWS
  VPWS-Id: 100, Name: vpws1 Status: UP, Mapped Tunnel: RSVP, ftn_ix 4, color 1
```

The following sample shows the detailed L3VPN service information of specific tunnel policy by name.

```
#show tunnel-policy tpl service l3vpn details
Total services using tunnel-policy tpl:
L3VPN
  VRF id 2
    FTN id: 1, Prefix: 20.10.30.0/24, Status: Up, Label: 25664, Mapped Tunnel: RSVP, ftn_ix 4, color
1
    FTN id: 2, Prefix: 3001::/64, Status: Up, Label: 25664, Mapped Tunnel: RSVP, ftn_ix 4, color 1
```

The following sample shows the detailed 6PE service information of specific tunnel policy by name.

```
#show tunnel-policy tpl service 6pe details
Total services using tunnel-policy tpl:
6PE
  FTN id: 3, Prefix: 3002::/64, Status: Up, Label: 25665, Mapped Tunnel: RSVP, ftn_ix 4, color 1
```

The following sample shows the detailed EVPN service information of specific tunnel policy by name.

```
#show tunnel-policy tpl service evpn details
Total services using tunnel-policy tpl:
EVPN
  EVPN-Id: 30, Tunnel count: 1
    Peer: 105.1.1.1, Mapped Tunnel: SR_POLICY, ftn_ix 3, color 202
  EVPN-Id: 1001, Tunnel count: 2
    Peer: 105.1.1.1, Mapped Tunnel: SR_POLICY, ftn_ix 3, color 202   Peer: 101.1.1.1, Mapped Tunnel:
SR_POLICY, ftn_ix 4, color 204
```

tunnel-policy

Use this command to create a tunnel policy.

Use the `no` parameter with this command to un-configure the tunnel policy. Tunnel policy can be un-configured only if no service is using it.

Command Syntax

```
tunnel-policy NAME  
no tunnel-policy NAME
```

Parameters

NAME

Name to be used for tunnel policy

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 4.2.

Example

```
#configure terminal  
(config)#tunnel-policy policy1
```

tunnel-select-policy (For L3VPN service)

Use this command to set tunnel-policy for a VRF instance. This will be used while selecting transport for L3VPN service.

Use the `no` parameter with this command to un-configure tunnel-policy from VRF instance.

Command Syntax

```
tunnel-select-policy NAME  
no tunnel-select-policy
```

Parameters

NAME

Name of tunnel-policy

Command Mode

IP VRF mode

Applicability

This command was introduced in OcNOS version 4.2.

Examples

```
#configure terminal  
(config)#ip vrf vrf1  
(config-vrf)#tunnel-select-policy policy1
```

MPLS LSP BFD Commands

This chapter describes each ICCP (Inter-Chassis Communication Protocol) command.

| | |
|-------------------------|------|
| lsp ping-interval | 1617 |
| min-tx | 1618 |
| mpls bfd | 1619 |
| shutdown | 1620 |

lsp ping-interval

Use this command to configure periodic MPLS ping echo request for BFD session over MPLS LSP.

Use `no` form of the command to set default periodic MPLS ping echo request interval.

Command Syntax

```
lsp-ping-interval <0-1000>  
no lsp-ping-interval
```

Parameters

<0-1000>

Ping interval in sec

Command Mode

MPLS BFD mode.

Applicability

This command was introduced before OcNOS version 5.0.

Example

```
#configure terminal  
(config)#mpls bfd ldp 1.1.1.1/32  
(config-mpls-bfd)# lsp-ping-interval 10
```

min-tx

Use this command to configure BFD transmit and receive intervals, and the hello multiplier value for MPLS LSP BFD session.

Use the `no` form of the command to set the intervals and multiplier to their default values.

Command Syntax

```
min-tx <50-999> min-rx <50-999> multiplier <3-50>  
no min-tx <50-999> min-rx <50-999> multiplier <3-50>
```

Parameters

min-tx

Set BFD Min Tx interval

<50-999>

Desired Tx interval in ms.

min-rx

Set BFD Min Rx interval

<50-999>

Desired Rx interval in ms.

Multiplier

Set BFD Detection Multiplier

<3-50>

Multiplier value

Defaults

The default for the transmit and receive intervals is 10 milliseconds. The default hello multiplier value is 3.

Command Mode

MPLS BFD mode.

Applicability

This command was introduced before OcNOS version 5.0.

Example

```
#configure terminal  
(config)#mpls bfd ldp 1.1.1.1/32  
(config-mpls-bfd)# min-tx 10 min-rx 10 multiplier 3
```

mpls bfd

Use this command to set up a BFD session for an MPLS LSP.

Use the `no` form of the command to remove the BFD session from the MPLS LSP.

Command Syntax

```
mpls bfd ((ldp A.B.C.D/M)|(rsvp tunnel-name NAME)|(static A.B.C.D/M))
mpls bfd static all
no mpls bfd ((ldp A.B.C.D/M)|(rsvp tunnel-name NAME)|(static A.B.C.D/M))
no mpls bfd static all
```

Parameters

ldp

Configures a BFD session for an LDP-signaled MPLS LSP.

ldp A.B.C.D/M

Specifies the LDP FEC (prefix) for the BFD session.

rsvp

Configures a BFD session for an RSVP-signaled MPLS LSP.

rsvp tunnel-name NAME

Specifies the RSVP tunnel name for the BFD session.

static A.B.C.D/M

Configures a BFD session for a specific static MPLS LSP.

static all

Configures BFD sessions for all static MPLS LSPs.

Command Mode

Configure mode

Applicability

Introduced before OcNOS version 5.0. Removed the command `mpls bfd (ldp|rsvp) all` in OcNOS version 6.6.1.

Example

```
#configure terminal
(config)#mpls bfd rsvp tunnel-name t1
(config-mpls-bfd)#exit

(config)#mpls bfd ldp 1.1.1.1/32
(config-mpls-bfd)#exit
```

shutdown

Use this command to configure MPLS LSP BFD session in Admin Down state or to disable BFD session.

Use `no` form of the CLI to enable BFD session.

Command Syntax

```
shutdown  
no shutdown
```

Parameters

None

Command Mode

MPLS BFD mode.

Applicability

This command was introduced before OcNOS version 5.0.

Example

```
#configure terminal  
(config)#mpls bfd ldp 1.1.1.1/32  
(config-mpls-bfd)# shutdown
```

L2VPN over L2 Subinterface Commands

This chapter is a reference for the L2VPN commands over L2 subinterface.

| | |
|------------------------------------|------|
| access-if-vpws | 1622 |
| mpls-l2-circuit | 1623 |
| vc-mode | 1624 |
| access-if-vpls | 1625 |
| mpls-vpls | 1626 |
| learning disable | 1627 |
| split-horizon (subinterface) | 1628 |
| static-mac | 1629 |

access-if-vpws

Use this command in L2 subinterface mode to create the vpws access-port.

Use the `no` form of this command to delete the vpws access port.

Command Syntax

```
access-if-vpws  
no access-if-vpws
```

Parameters

None

Command Mode

L2 Sub-interface mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal  
(config)#interface xel.1 switchport  
(config-if)#access-if-vpws  
(config-acc-if-vpws)#end
```

mpls-l2-circuit

Use this command to bind an mpls-l2-circuit to an L2 subinterface

Use the `no` form of this command to unbind an mpls-l2-circuit from an L2 subinterface

Command Syntax

```
mpls-l2-circuit NAME (primary | secondary)
no mpls-l2-circuit NAME
```

Parameters

NAME

A string identifying the MPLS Layer-2 Virtual Circuit

primary

Identify L2 circuit as the primary link

secondary

Identify L2 circuit as the secondary link; the secondary link is not activated unless the primary link fails

Command Mode

access-if-vpws mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
(config)#interface xe1.1 switchport
(config-if)#access-if-vpws
(config-acc-if-vpws)#mpls-l2-circuit vc1 primary
(config-acc-if-vpws)#mpls-l2-circuit vc2 secondary
(config-acc-if-vpws)#end
```

vc-mode

Use this command to configure vc-mode as revertive for an mpls l2-circuit.

Use the `no` form of this command to remove the vc-mode.

Command Syntax

```
vc-mode revertive timer  
no vc-mode revertive
```

Parameters

timer

(Optional) Specifies the duration, in seconds, that the system waits before reverting the Virtual Circuit (VC) to the primary path once it becomes available. The valid range is <0–3600> seconds and the default value is 0.

Command Mode

access-if-vpws mode or vpls-spoke mode

Applicability

This command was introduced in OcNOS version 5.0 and added an optional parameter `timer` in OcNOS version 7.0.0.

Examples

The following example illustrates how to configure VC revertive mode on an access interface for VPWS:

```
#configure terminal  
(config)#interface xel.1 switchport  
(config-if)#access-if-vpws  
(config-acc-if-vpws)#mpls-l2-circuit vc1 primary  
(config-acc-if-vpws)#mpls-l2-circuit vc2 secondary  
(config-acc-if-vpws)#vc-mode revertive  
(config-acc-if-vpws)#end
```

The following example illustrates how to configure VC revertive mode with a 20-second revertive timer for a VPLS instance:

```
#configure terminal  
(config)#mpls vpls vpls200 200  
(config-vpls)#vpls-vc vc2000  
(config-vpls-spoke)#secondary vc2001  
(config-vpls-spoke)#vc-mode revertive timer 20  
(config-vpls-spoke)#exit-spoke  
(config-vpls)#exit-vpls
```

access-if-vpls

Use this command in L2 subinterface mode to create the vpls access-port.

Use the `no` form of this command to delete the vpls access port

Command Syntax

```
access-if-vpls  
no access-if-vpls
```

Parameters

None

Command Mode

L2 Sub-interface mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal  
(config)#interface xel.1 switchport  
(config-if)#access-if-vpls  
(config-acc-if-vpls)#end
```

mpls-vpls

Use this command to bind a VPLS instance to an L2 subinterface.

Use the `no` form of this command to unbind vpls instance from a L2 subinterface.

Command Syntax

```
mpls-vpls NAME  
no mpls-vpls NAME
```

Parameters

NAME

VPLS instance name

Command Mode

access-if-vpls mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal  
(config)#interface xel.2 switchport  
(config-if)#access-if-vpls  
(config-acc-if-vpls)#mpls-vpls v100  
(config-acc-if-vpls)#end
```

learning disable

Use this command to disable learning on a particular Attachment Circuit (AC) interface.

Use the `no` form of this command to enable learning on a particular AC interface.



Note: This command disables MAC learning only on that interface.

Command Syntax

```
learning disable
no learning disable
```

Parameters

None

Command Mode

access-if-vpls mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal
(config)#interface xel.2 switchport
(config-if)#access-if-vpls
(config-acc-if-vpls)#learning disable
(config-acc-if-vpls)#end
```

split-horizon (subinterface)

Use this command to configure split-horizon group for L2 sub-interface.

Use the `no` form of this command to remove split-horizon group from L2 sub-interface.

Command Syntax

```
split-horizon group (network | access1 | access2)
no split-horizon group
```

Parameters

network

Split-horizon group name network

access1

Split-horizon group name access1

access2

Split-horizon group name access2

Default

By default, split horizon is disabled.

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 6.3.0 and applicable only on Qumran1 devices..

Examples

```
#configure terminal
(config)#interface xel.1 switchport
(config-if)#split-horizon group access1
(config-if)#no split-horizon group
```

static-mac

Use this command to add a static MAC address to an attachment circuit specific for a VPLS instance.

Use the `no` parameter with this command to remove static MAC address.

Command Syntax

```
static-mac XXXX.XXXX.XXXX  
no static-mac XXXX.XXXX.XXXX
```

Parameters

XXXX.XXXX.XXXX

MAC address in HHHH.HHHH.HHHH format

Command Mode

access-if-vpls mode

Applicability

This command was introduced in OcNOS version 5.0.

Examples

```
#configure terminal  
(config)#interface xel.2 switchport  
(config-if)#access-if-vpls  
(config-acc-if-vpls)#static-mac 0091.0000.1010  
(config-acc-if-vpls)#end
```

ETHERNET VIRTUAL PRIVATE NETWORK CONFIGURATION

| | |
|--|------|
| EVPN MPLS Configuration | 1633 |
| Overview | 1633 |
| VPN Terminology | 1633 |
| EVPN MPLS Single Homing | 1634 |
| EVPN MPLS Single Homing Over BGP-LU | 1646 |
| EVPN MPLS Multihoming | 1662 |
| EVPN MPLS Multihoming over BGP-LU | 1680 |
| EVPN-MPLS L2CP Tunneling | 1703 |
| EVPN-MPLS MAC Statistics | 1704 |
| EVPN-MPLS Control-Word Support | 1705 |
| EVPN-MPLS RSVP Multipath | 1708 |
| EVPN-ELINE CFM Single Homing | 1711 |
| LDP ECMP Configuration | 1722 |
| Topology | 1722 |
| Configurations | 1722 |
| EVPN Active-Standby | 1752 |
| Overview | 1752 |
| Prerequisites | 1754 |
| Configuration | 1755 |
| EVPN MPLS Active-Standby MultiHoming Configuration | 1755 |
| EVPN SR Active-Standby Multi-Homing Configuration | 1798 |
| Implementation Examples | 1852 |
| Revised CLI Commands | 1852 |
| CLI Commands | 1852 |
| Glossary | 1856 |
| EVPN MPLS E-Tree Scenario 1 | 1859 |
| Overview | 1859 |
| Prerequisites | 1860 |
| Configuration | 1863 |
| Implementation Examples | 1885 |
| E-Tree Scenario 1 Commands | 1886 |
| Troubleshooting | 1886 |
| Glossary | 1887 |
| EVPN MPLS E-Tree Scenario 2 | 1888 |
| Overview | 1888 |

| | |
|--|------|
| E-Tree Scenario 2 Prerequisites | 1890 |
| E-Tree Scenario 2 Configuration | 1893 |
| E-Tree Scenario 2 Implementation Examples | 1918 |
| E-Tree Scenario 2 Troubleshooting | 1918 |
| E-Tree Scenario 2 Revised Commands | 1919 |
| E-Tree Scenario 2 Glossary | 1921 |
| Flow-Aware Transport for EVPN over MPLS | 1924 |
| Overview | 1924 |
| FAT Prerequisites | 1924 |
| FAT Configuration | 1928 |
| FAT Implementation Example | 1934 |
| FAT Revised Command | 1935 |
| FAT Glossary | 1937 |
| EVPN MAC Limit Enforcement at EVI and AC for EVPN-MPLS | 1939 |
| Overview | 1939 |
| Feature Characteristics | 1939 |
| Benefits | 1939 |
| Prerequisites | 1939 |
| Limitations | 1940 |
| Topology | 1940 |
| Configuration | 1940 |
| Validation | 1941 |
| CLI Commands | 1942 |
| Anycast Gateway Routing for Multiple Subnets in EVPN-IRB | 1949 |
| Overview | 1949 |
| Feature Characteristics | 1949 |
| Benefits | 1949 |
| Topology | 1950 |
| Configuration | 1950 |
| Validation | 1963 |
| Abbreviations | 1981 |
| CFM over EVPN-MPLS for ELINE MultiHoming | 1982 |
| Overview | 1982 |
| Feature Characteristics | 1982 |
| Benefits | 1983 |
| Prerequisites | 1983 |
| Configuration | 1984 |
| Implementation Examples | 2000 |
| Troubleshooting | 2000 |
| Abbreviations | 2001 |

| | |
|--|------|
| Glossary | 2001 |
| EVPN EPL Link-Loss Forwarding | 2003 |
| Overview | 2003 |
| Feature Characteristics | 2003 |
| Configuration of LLF for an EVPN Ethernet Private Line (EPL) | 2003 |
| CLI Commands | 2005 |
| Glossary | 2006 |
| TWAMP over EVPN Configuration | 2007 |
| Topology | 2007 |
| Configure TWAMP over EVPN for LDP | 2008 |
| Validation | 2013 |

EVPN MPLS Configuration

This chapter includes step-by-step configurations for EVPN MPLS for Single Homing and Multi Homing.

Overview

Ethernet VPN (EVPN) solution provides Ethernet multipoint services over MPLS networks. EVPN operates in contrast to the existing Virtual Private LAN Service (VPLS) by enabling control-plane based MAC learning. PEs participating in the EVPN instances learn customer MAC routes in control-plane using MP-BGP protocol. Control-plane MAC learning brings a number of benefits that allow EVPN to address the VPLS shortcomings, including support for multi-homing with per-flow load balancing.

In EVPN, PEs advertise the MAC addresses learned from the CEs that are connected to them, along with an MPLS label, to other PEs in the control plane using Multiprotocol BGP (MP-BGP). Control-plane learning enables load balancing of traffic to and from CEs that are multihomed to multiple PEs. This is in addition to load balancing across the MPLS core via multiple LSPs between the same pair of PEs. It also improves convergence times in the event of certain network failures.



Note: The EVPN is supported over static LSP, LDP, BGP-LU transports.

VPN Terminology

MAC-VRF: A virtual routing and forwarding table for storing MACs on a PE for specific bridge domain.

CE: Customer Edge device, e.g., a host, router, or switch.

PE: Provider edge device

EVI: An EVPN instance spanning the Provider Edge (PE) devices participating in that EVPN.



Note: Ethernet Segment (ES): Set of Ethernet links connected between CE and PE. Single CE can be connected to multiple PEs.

Ethernet Segment Identifier (ESI): A unique non-zero identifier that identifies an Ethernet segment is called an ESI

Ethernet Tag: An Ethernet tag identifies a particular broadcast domain, e.g., a VLAN. An EVPN instance consists of one or more broadcast domains.

Benefits

The EVPN control-plane MAC learning has the following benefits:

- Eliminate flood and learn mechanism as hosts are learned over control plane.
- OcNOS supports both dynamically learned hosts and statically configured hosts, which are advertised/learned over the EVPN control plane.
- Fast-reroute, resiliency, and faster convergence in case of multihoming
- Load balancing of traffic to and from CEs that are multihomed to multiple PE's.

The following EVPN types are supported:

- Single-homed CE: One CE is connected to One PE device.
- Multihomed CE: One CE is connected to Multiple PE devices. OcNOS supports dual-homed CEs with all-active multi homing mode.

Route Types

These EVPN route types are supported:

- Route Type 1: Ethernet Auto-Discovery (AD) Route

The Ethernet (AD) routes are advertised on per EVI and per ESI basis. These routes are sent per ES. They carry the list of EVIs that belong to the ES.

This route is advertised when multihomed CEs already exist.

- Route Type 2: MAC/IP Advertisement Route

The host's IP and MAC addresses are advertised to the peers within NLRI. The control plane learning of MAC addresses reduces unknown unicast flooding.

- Route Type 3: Inclusive Multicast Ethernet Tag Route

This route establishes the connection for broadcast, unknown unicast, and multicast (BUM) traffic from a source PE to a remote PE.

This route is advertised on per VLAN and per ESI basis.

- Route Type 4: Ethernet Segment Route

Ethernet segment routes enable to connect a CE device to two or PE devices.

Ethernet segment routes enable the discovery of connected PE devices that are connected to the same Ethernet segment.

EVPN auto route target will be supported under MAC VRF.

In EVPN-VPWS the auto-discovery of peer PE nodes is done with the pair of Ethernet A-D routes. Inclusive Multicast route does not have participation on auto-discovery unlike ELAN-EVPN VPLS. Since there is no mac-advertisement, MAC-IP route is not applicable




Note: Only Sub-interface supported as Access-port for EVPN-MPLS.

EVPN MPLS Single Homing

Topology

The following figure depicts the Single Homed topology for the EVPN MPLS configuration examples for both ELINE and ELAN service with LDP as underlay MPLS path.


Figure 48. EVPN MPLS Single Homing configuration**PE1: Loopback Interface**

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.1/32 secondary | Configure IP address on loopback interface. |
| (config-if)#commit | Commit the transaction. |
| (config-if)#exit | Exit interface mode |

PE1: Global EVPN MPLS Command

| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#evpn mpls vtep-ip-global 10.143.73.1 | Configuring VTEP global IP to loopback IP |
| (config)#commit | Commit candidate configuration to be running configuration Note: Reload is required after Enabling/Disabling EVPN MPLS Feature. |

PE1: Global LDP

| | |
|---|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.1 | Set the router ID to IP address 10.143.73.1 |
| (config-router)#transport-address ipv4 10.143.73.1 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| <div style="border: 1px solid black; padding: 5px;">  Note: It is preferable to use the loopback address as the transport address. </div> | |
| (config-router)#targeted-peer ipv4 10.143.73.3 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit the transaction. |

PE1: Interface Configuration Network Side

| | |
|--|--|
| (config)#interface xe3 | Enter the Interface mode for xe3. |
| (config-if)#ip address 10.255.128.2/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |



Note: For RSVP Configuration refer [RSVP-TE Configuration \(page 643\)](#)

PE1: OSPF Configuration

| | |
|--|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.1 | Router-ID configurations |
| (config-router)#network 10.143.73.1/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)#network 10.255.128.2/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit the transaction. |

PE1: BGP Configuration

| | |
|---|---|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.3 remote-as 65010 | Configuring PE3 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.3 update-source lo | Source of routing updates as loopback |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.3 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config-router)#commit | Commit the transaction. |

PE1: MAC VRF Configuration

| | |
|------------------------------------|---|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.1:2 | Configuring Route-Distinguisher value 10.143.73.1:2 |
| (config-vrf)#route-target both 2:2 | Configuring import and export value as 2:2 |
| (config-vrf)#exit | Exiting VRF Mode |

PE1: MAC VRF Configuration with Auto route target

| | |
|---|--|
| (config)#mac vrf vpls1001 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.1:1001 | Configuring Route-Distinguisher value 10.143.73.1:1001 |
| (config-vrf)#route-target both evpn-auto-rt | Configuring import and export value as evpn-auto-rt. Route target will be derived automatically. Support: route-target export route-target import |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#commit | Commit the transaction. |

PE1: EVPN-VPWS (ELINE) and MAC VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 2 xconnect target-mpls-id 252 | Configure the EVPN-VPWS identifier with source identifier 2 and target identifier 252 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE1: EVPN-VPWS (ELINE) Access Port Configuration

| | |
|--|---|
| (config)#interface xe6 | Enter the Interface mode for xe6. |
| (config-if)#interface xe6.2 switchport | Creating L2 sub interface of physical interface xe6. |
| (config-if)#encapsulation dot1q 2 | Setting Encapsulation to dot1q with VLAN ID 2. Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if)#map vpn-id 2 | Map vpn-id 2 to interface xe6.2 (VPWS). |
| (config-access-if)#exit | Exiting out of access interface mode |

PE1: EVPN-VPLS (ELAN) and MAC VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 1001 | Configure the EVPN-VPLS identifier with identifier 1001. |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping VRF "vpls1001" to EVPN-VPLS identifier. |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |


PE1: EVPN-VPLS (ELAN) Access Port Configuration

| | |
|---|---|
| (config-if)#interface xe6.1001 switchport | Creating L2 sub interface of physical interface xe6 |
| (config-if)#encapsulation dot1q 1001 | Setting Encapsulation to dot1q with VLAN ID 1001 Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if)#map vpn-id 1001 | Map vpn-id 1001 to interface xe6.1001 (VPLS) |
| OcNOS (config-access-if)#commit | Commit candidate configuration to be running configuration |

P: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.2/32 secondary | Configure IP address on loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |

P: Global LDP

| | |
|--|---|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.2 | Set the router ID to IP address 10.255.128.25 |
| (config-router)#transport-address ipv4 10.143.73.2 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit the transaction |

P: Interface Configuration

| | |
|--|--|
| (config)#interface xe3 | Enter the Interface mode for xe3. |
| (config-if)#ip address 10.255.128.3/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter the Interface mode for xe5 |

| | |
|---|--|
| (config-if)#ip address 10.255.128.22/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |


P: OSPF Configuration

| | |
|---|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.2 | Setting the Router ID as Loopback IP |
| (config-router)#network 10.143.73.2/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)#network 10.255.128.2/31 area 0.0.0.0 | Advertise network address in OSPF that comes under same subnet. |
| (config-router)#network 10.255.128.22/31 area 0.0.0.0 | Advertise xe5 network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.3/32 secondary | Configure IP address on loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |

PE2: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.3 | Set the router ID to IP address 10.143.73.3 |
| (config-router)#transport-address ipv4 10.143.73.3 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| | <div>  Note: It is preferable to use the loopback address as the transport address. </div> |
| (config-router)#targeted-peer ipv4 10.143.73.1 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit router mode |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit the transaction. |

PE2: Global EVPN MPLS Command

| | |
|---|---|
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#commit | Commit candidate configuration to be running configuration Note: Reload is required after Enabling/Disabling EVPN MPLS Feature |
| (config)#evpn mpls vtep-ip-global 10.143.73.3 | Configuring VTEP global IP to loopback IP |
| (config)#commit | Commit the transaction. |

Interface Configuration Network Side:

| | |
|---|--|
| (config)#interface xe5 | Enter the Interface mode for xe5. |
| (config-if)#ip address 10.255.128.23/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |

PE2: OSPF Configuration

| | |
|---|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.3 | Router-ID configurations |
| (config-router)#network 10.143.73.3/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)#network 10.255.128.22/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit the transaction. |

PE2: BGP Configuration

| | |
|---|---|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.1 remote-as 65010 | Configuring PE1 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.1 update-source lo | Source of routing updates as loopback |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.1 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config)#commit | Commit the transaction. |

PE2: MAC VRF Configuration

| | |
|------------------------------------|---|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.3:2 | Configuring Route-Distinguisher value 10.143.73.3:2 |
| (config-vrf)#route-target both 2:2 | Configuring import and export value as 2:2 Support: route-target export route-target import |
| (config-vrf)#exit | Exiting VRF Mode |

PE2: MAC VRF Configuration with Auto route target

| | |
|---|---|
| (config)#mac vrf vpls1001 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.1:1001 | Configuring Route-Distinguisher value 10.143.73.1:1001 |
| (config-vrf)#route-target both evpn-auto-rt | Configuring import and export value as evpn-auto-rt. Route target will be derived automatically. |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#commit | Commit the transaction. |

PE2: EVPN-VPWS (ELINE) and VRF Mapping

| | |
|---|--|
| (config)#evpn mpls id 252 xconnect target-mpls-id 2 | Configure the EVPN-ELINE identifier with source identifier 252 and target identifier 2 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-ELINE identifier |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: EVPN-VPWS (ELINE) Access Port Configuration

| | |
|---|---|
| (config)#interface xe2 | Enter the Interface mode for xe2. |
| (config-if)#description access-side-int | Giving Interface Description. |
| (config-if)#interface xe2.2 switchport | Creating L2 sub interface of physical interface xe2. |
| (config-if)#encapsulation dot1q 2 | Setting Encapsulation to dot1q with VLAN ID 2 Supported Encapsulation: dot1ad, dot1q, untagged, default. |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if)#map vpn-id 252 | Map vpn-id 252 to interface xe2.2 (VPWS) |
| (config-access-if)#exit | Exiting out of access interface mode |

PE2: EVPN-VPLS (ELAN) and VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 1001 | Configure the EVPN-ELAN identifier with identifier 1001. |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping VRF "vpls1001" to EVPN-ELAN identifier. |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: EVPN -VPLS (ELAN) Access Port Configuration

| | |
|---|---|
| (config-if)#interface xe2.1001 switchport | Creating L2 sub interface of physical interface xe2 |
| (config-if)#encapsulation dot1q 1001 | Setting Encapsulation to dot1q with VLAN ID 1001 Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if)#map vpn-id 1001 | Map vpn-id 1001 to interface xe2.1001 (VPLS) |
| (config)#commit | Commit candidate configuration to be running configuration |

Validation**PE1: E-LAN or EVPN-VPLS**

```

PE1#sh mac vrf vpls1001
VRF vpls1001, FIB ID 4098
Router ID is not set
Interfaces:
VRF vpls1001; default RD 10.143.73.3:1001
Evpn Auto RT:100:1073742025
Import VPN route-target communities
Evpn Auto RT:100:1073742025
No import route-map
No export route-map
VPNv4 label allocation mode: per-vrf
VPNv6 label allocation mode: per-vrf
import-vnid: 1001
export-vnid: 1001

```

```

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status          Up/Down         Update          evpn-id
=====
10.143.73.1      10.143.73.3     Installed       00:01:03        00:01:03        1001

```

Total number of entries are 1

```

PE1#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels

```

| Destination Label | Status | evpn-id | Network-Intf | Tunnel-Label | Local | | Remote | |
|----------------------|--------|---------|--------------|--------------|----------|----------|----------|----------|
| | | | | | MC-Label | UC-Label | MC-Label | UC-Label |

```

=====
10.143.73.3      Installed    1001      xe3        24320      640      17      640      17

Total number of entries are 1

PE1#show evpn mpls id 1001
EVPN-MPLS Information
=====
Codes: NW - Network Port
      AC - Access Port
      (u) - Untagged

VPN-ID   EVI-Name      EVI-Type Type Interface ESI              VLAN DF-Status Src-
Addr          Dst-Addr

-----
1001     vpls1001      L2      NW    ----      ----              ----  ----
      10.143.73.1    10.143.73.3
1001     vpls1001      --      AC    xe6.1001  --- Single Homed Port ---  ----  ----
      ----

Total number of entries are 2

PE1#show evpn mpls mac-table
=====
=====
EVPN MPLS MAC Entries
=====
=====
VNID      Interface VlanId Inner-VlanId Mac-Addr      VTEP-
Ip/ESI              Type      Status      AccessPortDesc
-----
1001      ----      ----      ----      00ff.2300.0000 10.143.73.3      Dynamic
Remote  -----      -----
1001      ----      ----      ----      0023.1001.0001 10.143.73.3      Dynamic
Remote  -----      -----
1001     xe6.1001      ----      ----      0045.1001.0001 10.143.73.1      Dynamic
Local   -----      -----
1001     xe6.1001      ----      ----      0071.1001.0001 10.143.73.1      Dynamic
Local   -----      -----

```

PE1: E-LINE or EVPN-VPWS

```

PE1#show evpn mpls xconnect
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local      Remote      Connection-Details
=====
=====
VPN-ID      EVI-Name      MTU  VPN-ID      Source      Destination      PE-
IP          MTU  Type  NW-Status
=====
2           ----      1500  252         xe6.2        --- Single Homed Port ---
      10.143.73.3    1500  AC-NW  NW-SET

Total number of entries are 1

```

```

PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====
=
10.143.73.1      10.143.73.3      Installed    00:01:10      00:01:10      2            252

Total number of entries are 1

PE1#show evpn mpls xconnect tunnel label
EVPN-MPLS Network tunnel labels
              Local      Remote
Remote
Destination  Status      VPWS-ID  VPWS-ID  Network-Intf  Tunnel-Label  MC-Label  UC-Label  MC-
Label  UC-Label
=====
=====
10.143.73.3    Installed    2         252        xe3          24320         --         16         -
-         16

Total number of entries are 1

PE1#show evpn mpls xconnect id 2
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local          Remote      Connection-Details
=====
=====
VPN-ID      EVI-Name      MTU  VPN-ID      Source      Destination      PE-
IP          MTU  Type  NW-Status
=====
=====
2           ----          1500  252         xe6.2        --- Single Homed Port ---
10.143.73.3  1500  AC-NW  NW-SET

Total number of entries are 1

```

PE2: ELAN or EVPN-VPLS

```

PE2#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      evpn-id
=====
=====
10.143.73.3      10.143.73.1      Installed    00:04:03      00:04:03      1001

Total number of entries are 1

PE2#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels
              Local      Remote
Destination  Status      evpn-id  Network-Intf  Tunnel-Label  MC-Label  UC-Label  MC-Label  UC-
Label
=====
=====
10.143.73.1    Installed    1001      xe5          24321         640         17         640         --

Total number of entries are 1

```



```
PE2#show evpn mpls id 1001
```

```
EVPN-MPLS Information
```

```
=====
```

```
Codes: NW - Network Port
```

```
AC - Access Port
```

```
(u) - Untagged
```

| VPN-ID Addr | EVI-Name Dst-Addr | EVI-Type | Type | Interface | ESI | VLAN | DF-Status | Src- |
|----------------|-------------------------|----------|------|-----------|------|-------------------|-----------|------|
| 1001 | vpls1001 10.143.73.3 | L2 | NW | ---- | ---- | ---- | ---- | |
| 1001 | vpls1001 ---- | -- | AC | xe2.1001 | --- | Single Homed Port | --- | ---- |

```
Total number of entries are 2
```

```
PE2#show evpn mpls mac-table
```

```
=====
```

```
EVPN MPLS MAC Entries
```

```
=====
```

| VNID Ip/ESI | Interface | VlanId | Inner-VlanId | Mac-Addr Status | VTEP- AccessPortDesc | |
|----------------|-----------|--------|--------------|--------------------|-------------------------|---------|
| 1001 | ---- | ---- | ---- | 0045.1001.0001 | 10.143.73.1 | Dynamic |
| Remote | ----- | ----- | | | | |
| 1001 | ---- | ---- | ---- | 0071.1001.0001 | 10.143.73.1 | Dynamic |
| Remote | ----- | ----- | | | | |
| 1001 | xe2.1001 | ---- | ---- | 00ff.2300.0000 | 10.143.73.3 | Dynamic |
| Local | ----- | ----- | | | | |
| 1001 | xe2.1001 | ---- | ---- | 0023.1001.0001 | 10.143.73.3 | Dynamic |
| Local | ----- | ----- | | | | |

PE2: ELINE or EVPN-VPWS

```
PE2#show evpn mpls xconnect
```

```
EVPN-MPLS Xconnect Info
```

```
=====
```

```
AC-AC: Local-Cross-connect
```

```
AC-NW: Cross-connect to Network
```

```
AC-UP: Access-port is up
```

```
AC-DN: Access-port is down
```

```
NW-UP: Network is up
```

```
NW-DN: Network is down
```

```
NW-SET: Network and AC both are up
```

| Local | | Remote | Connection-Details |
|-------------|----------|-----------|--------------------|
| ===== | | ===== | |
| VPN-ID | EVI-Name | MTU | VPN-ID |
| IP | MTU Type | NW-Status | Source |
| ===== | | ===== | Destination |
| ===== | | ===== | PE- |
| 252 | ---- | 1500 | 2 |
| 10.143.73.1 | 1500 | AC-NW | NW-SET |
| | | | xe2.2 |
| | | | --- |
| | | | Single Homed Port |
| | | | --- |

```
Total number of entries are 1
```

```
PE2#show evpn mpls xconnect tunnel
```

```
EVPN-MPLS Network tunnel Entries
```

| Source | Destination | Status | Up/Down | Update | local-evpn-id | remote- |
|---------|-------------|--------|---------|--------|---------------|---------|
| evpn-id | | | | | | |

```
=====
===
10.143.73.3      10.143.73.1      Installed      00:04:10      00:04:10      252      2

Total number of entries are 1

PE2#show evpn mpls xconnect tunnel label
EVPN-MPLS Network tunnel labels

      Local      Remote      Local
Remote
Destination      Status      VPWS-ID      VPWS-ID      Network-Intf      Tunnel-Label      MC-Label      UC-Label      MC-
Label      UC-Label
=====
10.143.73.1      Installed      252      2      xe5      24321      --      16      -
-      16

Total number of entries are 1

PE2#show evpn mpls xconnect id 252
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local      Remote      Connection-Details
=====
VPN-ID      EVI-Name      MTU      VPN-ID      Source      Destination      PE-
IP      MTU      Type      NW-Status
=====
252      ----      1500      2      xe2.2      --- Single Homed Port ---
10.143.73.1      1500      AC-NW      NW-SET


Total number of entries are 1
```

EVPN MPLS Single Homing Over BGP-LU

Topology

depicts the Single Homed topology for the EVPN MPLS configuration examples for both ELINE and ELAN service with BGP-LU as underlay MPLS path which in turn over another underlay of LDP/RSVP/SR MPLS paths.

Figure 49. EVPN MPLS Single Homing over LU configuration



Configuration**PE1: Loopback Interface**

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.1/32 secondary | Configure IP address on loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE1: Global EVPN MPLS Command

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#evpn mpls vtep-ip-global 10.143.73.1 | Configuring VTEP global IP to loopback IP |
| (config)#commit | Commit candidate configuration to be running configuration |

PE1: Global LDP

| | |
|---|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.1 | Set the router ID to IP address 10.143.73.1 |
| (config-router)# transport-address ipv4 10.143.73.1 | Configure transport address at LDP mode |
| (config-router)#targeted-peer ipv4 10.143.73.5 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE1: Interface Configuration Network Side

| | |
|--|--|
| (config)#interface xe3 | Enter the Interface mode for eth2. |
| (config-if)#ip address 10.255.128.8/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |



Note: For RSVP Configuration refer [RSVP-TE Configuration \(page 643\)](#).

PE1: OSPF Configuration

| | |
|--|-------------------------------------|
| (config)#router ospf 1 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.1 | Router-ID configurations |
| (config-router)#network 10.143.73.1/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)#network 10.255.128.8/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)#commit | Commit the transaction. |

PE1: BGP Configuration

| | |
|---|--|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config)#allocate-label all | Configure allocate all under router BGP |
| (config-router)#neighbor 10.143.73.5 remote-as 65010 | Configuring PE2 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.2 remote-as 65010 | Configuring ABR1 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.5 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.2 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 unicast | Entering into address family ipv4 unicast |
| (config-router-af)# network 10.143.73.1/32 | Advertise 10.143.73.1/31 network |
| (config-router-af)# exit-address-family | Exit from IPv4 address family |
| (config-router)# address-family ipv4 labeled-unicast | Entering into address family ipv4 labeled-unicast |
| (config-router-af)#neighbor 10.143.73.2 activate | Activate the IPv4 labeled-unicast neighbor |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.5 activate | Enabling EVPN Address family for neighbor |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE1: MAC VRF Configuration

| | |
|------------------------------------|---|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.1:2 | Configuring Route-Distinguisher value 10.143.73.1:2 |
| (config-vrf)#route-target both 2:2 | Configuring import and export value as 2:2 Support: route-target export route-target import |

| | |
|--|---|
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#mac vrf vpls1001 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.1:1001 | Configuring Route-Distinguisher value 10.143.73.1:1001 |
| (config-vrf)#route-target both 1001:1001 | Configuring import and export value as 1001:1001 Support: route-target export Support: route-target export |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE1: EVPN-VPWS (ELINE) and MAC VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 2 xconnect target-mpls-id 252 | Configure the EVPN-VPWS identifier with source identifier 2 and target identifier 252 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE1: EVPN -VPWS (ELINE) Access Port Configuration

| | |
|---|---|
| (config)#interface xe6 | Enter the Interface mode for xe6. |
| (config-if)#description access-side-int | Giving Interface Description. |
| (config-if)#interface xe6.2 switchport | Creating L2 sub interface of physical interface xe6. |
| (config-if)#encapsulation dot1q 2 | Setting Encapsulation to dot1q with VLAN ID 2 Supported Encapsulation: dot1ad, dot1q, untagged, default. |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if)#map vpn-id 2 | Map VPN-ID 2 to interface xe6.2 (VPWS). |
| (config-access-if)#exit | Exiting out of access interface mode. |

PE1: EVPN-VPLS (ELAN) and MAC VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 1001 | Configure the EVPN-VPLS identifier with identifier 1001 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-VPLS identifier |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE1: EVPN-VPLS (ELAN) Access Port Configuration

| | |
|---|---|
| (config-if)#interface xe6.1001 switchport | Creating L2 sub interface of physical interface xe6 |
| (config-if)#encapsulation dot1q 1001 | Setting Encapsulation to dot1q with VLAN ID 1001 Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if)#map vpn-id 1001 | Map vpn-id 1001 to interface xe6.1001 (VPLS) |
| (config-access-if)#commit | Commit candidate configuration to be running configuration |

ABR1: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.2/32 secondary | Configure IP address on loopback interface. |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

ABR1: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#transport-address ipv4 10.143.73.2 | Configure transport address under router ldp |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

ABR1: Interface Configuration

| | |
|---|--|
| (config)#interface xe3 | Enter the Interface mode for xe3. |
| (config-if)#ip address 10.255.128.9/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)#ip address 10.255.128.25/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |

| | |
|-----------------------------------|--|
| <code>(config-router)#exit</code> | Exit from router target peer and LDP mode |
| <code>(config)#commit</code> | Commit candidate configuration to be running configuration |

ABR1: OSPF Configuration

| | |
|--|--|
| <code>(config)#enable ext-ospf-multi-inst</code> | Enable multiple-instance capabilit |
| <code>(config)#router ospf 1</code> | Enter the Router OSPF mode. |
| <code>(config-router)#ospf router-id 10.143.73.2</code> | Setting the Router ID as Loopback IP |
| <code>(config-router)#network 10.143.73.2/32 area 0.0.0.0 instance-id 100</code> | Advertise loopback address in OSPF. |
| <code>(config-router)#network 10.255.128.25/31 area 0.0.0.0</code> | Advertise network address in OSPF |
| <code>(config)#router ospf 100</code> | Enter the Router OSPF mode. |
| <code>(config-router)#network 10.143.73.2/32 area 0.0.0.0 instance-id 1</code> | Advertise loopback address in OSPF |
| <code>(config-router)#network 10.255.128.9/31 area 0.0.0.0</code> | Advertise network address in OSPF that comes under same subnet |
| <code>(config-router)#exit</code> | Exit Router OSPF mode and return to Configure mode |
| <code>(config)#commit</code> | Commit candidate configuration to be running configuration |

ABR1: BGP-LU Configuration

| | |
|--|--|
| <code>(config)#router bgp 65010</code> | Enter the Router BGP mode, ASN: 65010 |
| <code>(config)#allocate-label all</code> | Configure allocate all under router BGP |
| <code>(config-router)#neighbor 10.143.73.1 remote-as 65010</code> | Configuring PE1 as iBGP neighbor using it's loopback IP |
| <code>(config-router)#neighbor 10.143.73.4 remote-as 65010</code> | Configuring ABR2 as iBGP neighbor using it's loopback IP |
| <code>(config-router)#neighbor 10.143.73.1 update-source lo</code> | Source of routing updates as loopback |
| <code>(config-router)#neighbor 10.143.73.4 update-source lo</code> | Source of routing updates as loopback |
| <code>(config-router)# address-family ipv4 unicast</code> | Entering into address family ipv4 unicast |
| <code>(config-router-af)# network 10.143.73.2/32</code> | Advertise the network |
| <code>(config-router-af)# exit-address-family</code> | Exit from IPv4 address family |
| <code>(config-router)# address-family ipv4 labeled-unicast</code> | Entering into address family ipv4 labeled-unicast |
| <code>(config-router-af)#neighbor 10.143.73.4 activate</code> | Activate the ipv4 labeled-unicast neighbor |
| <code>(config-router-af)#neighbor 10.143.73.4 next-hop-self</code> | Configure next-hopself for the ipv4 labeled-unicast neighbor |

| | |
|---|--|
| config-router-af)#neighbor 10.143.73.4 route-reflector-client | Configure neighbor as route reflector client |
| (config-router-af)#neighbor 10.143.73.1 activate | Activate the ipv4 labeled-unicast neighbor |
| config-router-af)#neighbor 10.143.73.1 route-reflector-client | Configure neighbor as route reflector client |
| (config-router-af)#neighbor 10.143.73.1 next-hop-self | Configure next-hopself for the ipv4 labeled-unicast neighbor |
| (config)#commit | Commit candidate configuration to be running configuration |

P2: Loopback Interface

| | |
|---|---|
| #configure terminal | Enter configuration mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.3/32 secondary | Configure IP address on loopback interface |
| (config-f)#commit | Commit the transaction |

P2: Global LDP

| | |
|------------------------|---------------------------|
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#commit | Commit the transaction |

P2: Interface Configuration

| | |
|---|--|
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)#ip address 10.255.128.24/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe6 | Enter the Interface mode for xe6 |
| (config-if)#ip address 10.255.128.35/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#commit | Commit the transaction |

P2: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.3 | Setting the Router ID as Loopback IP |
| (config-router)#network 10.143.73.3/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.24/31 area | Advertise network address in OSPF that comes |

| | |
|---|--|
| 0.0.0.0 | under same subnet |
| (config-router)#network 10.255.128.35/31 area 0.0.0.0 | Advertise xe5 network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit candidate configuration to be running configuration |

ABR2: Loopback Interface

| | |
|---|---|
| #configure terminal | Enter configuration mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.4/32 secondary | Configure IP address on loopback interface |
| (config-if)#commit | Commit the transaction |

ABR2: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#transport-address ipv4 10.143.73.4 | Configure transport address under router ldp |
| (config)#commit | Commit the transaction |

ABR2: Interface Configuration

| | |
|---|--|
| (config)#interface xe6 | Enter the Interface mode for xe6. |
| (config-if)#ip address 10.255.128.34/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe3 | Enter the Interface mode for xe3 |
| (config-if)#ip address 10.255.128.45/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#commit | Commit the transaction. |

ABR2: OSPF Configuration

| | |
|---|--|
| (config)#enable ext-ospf-multi-inst | Enable multiple-instance capabilit |
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.4 | Setting the Router ID as Loopback IP |
| (config-router)#network 10.143.73.4/32 area 0.0.0.0 instance-id 200 | Advertise loopback address in OSPF. |
| (config-router)#network 10.255.128.34/31 area | Advertise network address in OSPF that comes |

| | |
|---|--|
| 0.0.0.0 | under same subnet. |
| (config)#router ospf 200 | Enter the Router OSPF mode. |
| (config-router)#network 10.255.128.45/31 area 0.0.0.0 instance-id 100 | Advertise xe5 network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit candidate configuration to be running configuration |

ABR2: BGP-LU Configuration

| | |
|---|---|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config)#allocate-label all | Configure allocate all under router BGP |
| (config-router)#neighbor 10.143.73.2 remote-as 65010 | Configuring ABR1 as iBGP neighbor using its loopback IP |
| (config-router)#neighbor 10.143.73.5 remote-as 65010 | Configuring PE2 as iBGP neighbor using its loopback IP |
| (config-router)#neighbor 10.143.73.2 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.5 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 unicast | Entering into address family ipv4 unicast |
| (config-router-af)# network 10.143.73.4/32 | Advertise the network |
| (config-router-af)# exit-address-family | Exit from IPv4 address family |
| (config-router)# address-family ipv4 labeled-unicast | Entering into address family IPv4 labeled-unicast |
| (config-router-af)#neighbor 10.143.73.2 activate | Activate the IPv4 labeled-unicast neighbor |
| (config-router-af)#neighbor 10.143.73.2 next-hop-self | Configure next-hop-self for the IPv4 labeled-unicast neighbor |
| config-router-af)#neighbor 10.143.73.2 route-reflector-client | Configure neighbor as route reflector client |
| (config-router-af)#neighbor 10.143.73.5 activate | Activate the IPv4 labeled-unicast neighbor |
| config-router-af)#neighbor 10.143.73.5 route-reflector-client | Configure neighbor as route reflector client |
| (config-router-af)#neighbor 10.143.73.5 next-hop-self | Configure next-hop-self for the ipv4 labeled-unicast neighbor |
| (config-router-af)#commit | Commit the transaction. |

PE2: Loopback Interface


| | |
|---|---|
| #configure terminal | Enter configuration mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.5/32 secondary | Configure IP address on loopback interface |

| | |
|------------------|------------------------|
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |

PE2: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.5 | Set the router ID to IP address 10.143.73.5 |
| (config-router)#transport-address ipv4 10.143.73.5 | Configure transport address under router ldp |
| (config-router)#targeted-peer ipv4 10.143.73.3 | Configure targeted peer |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#commit | Commit the transaction |

PE2: Global EVPN MPLS Command

| | |
|---|--|
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#commit | Commit candidate configuration to be running configuration |
| |  Note: Reload is required after enabling or disabling EVPN MPLS. |
| (config)#evpn mpls vtep-ip-global 10.143.73.5 | Configuring VTEP global IP to loopback IP |
| (config)#commit | Commit the transaction |

PE2: Interface Configuration Network Side

| | |
|---|--|
| (config)#interface xe3 | Enter the Interface mode for xe3. |
| (config-if)#ip address 10.255.128.44/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-router)#commit | Commit the transaction. |

PE2: OSPF Configuration

| | |
|---|-------------------------------------|
| (config)#router ospf 200 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.5 | Router-ID configurations |
| (config-router)#network 10.143.73.5/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)#network 10.255.128.44/31 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)#commit | Commit the transaction. |

PE2: BGP Configuration

| | |
|---|---|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config)#allocate-label all | Configure allocate all under router BGP |
| (config-router)#neighbor 10.143.73.1 remote-as 65010 | Configuring PE1 as iBGP neighbor using its loopback IP |
| (config-router)#neighbor 10.143.73.4 remote-as 65010 | Configuring ABR2 as iBGP neighbor using its loopback IP |
| (config-router)#neighbor 10.143.73.1 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.4 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 unicast | Entering into address family ipv4 unicast |
| (config-router-af)# network 10.143.73.5/32 | Advertise the network |
| (config-router-af)# exit-address-family | Exit from IPv4 address family |
| (config-router)# address-family ipv4 labeled-unicast | Entering into address family ipv4 labeled-unicast |
| (config-router-af)#neighbor 10.143.73.4 activate | Activate the ipv4 labeled-unicast neighbor |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.1 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config-router)#commit | Commit the transaction |

PE2: MAC VRF Configuration

| | |
|--|--|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.5:2 | Configuring Route-Distinguisher value 10.143.73.3:2 |
| (config-vrf)#route-target both 2:2 | Configuring import and export value as 2:2 Support: route-target export route-target import |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#mac vrf vpls1001 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.5:1001 | Configuring Route-Distinguisher value 10.143.73.3:1001 |
| (config-vrf)#route-target both 1001:1001 | Configuring import and export value as 1001:1001 |
| (config-vrf)#commit | Commit the transaction |

PE2: EVPN-ELINE (VPWS) and VRF Mapping

| | |
|---|--|
| (config)#evpn mpls id 252 xconnect target-mpls-id 2 | Configure the EVPN-ELINE identifier with source identifier 252 and target identifier 2 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-ELINE identifier |

| | |
|----------------------------|---|
| (config-evpn-mpls) #commit | Commit the transaction. |
| (config-evpn-mpls) #exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: EVPN-VPWS (ELINE) Access Port Configuration

| | |
|--|---|
| (config) #interface xe2 | Enter the Interface mode for xe2. |
| (config-if) #description access-side-int | Giving Interface Description. |
| (config-if) #interface xe2.2 switchport | Creating L2 sub interface of physical interface xe2. |
| (config-if) #encapsulation dot1q 2 | Setting Encapsulation to dot1q with VLAN ID 2. Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if) #access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if) #map vpn-id 252 | Map VPN-ID 252 to interface xe2.2 (VPWS). |

PE2: EVPN-VPLS (ELAN) and VRF Mapping

| | |
|--|---|
| (config) #evpn mpls id 1001 | Configure the EVPN-ELAN identifier with identifier 1001 |
| (config-evpn-mpls) #host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-ELAN identifier |
| (config-evpn-mpls) #commit | Commit the transaction. |
| (config-evpn-mpls) #exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: EVPN-VPLS (ELAN) Access Port Configuration

| | |
|--|---|
| (config-if) #interface xe2.1001 switchport | Creating L2 sub interface of physical interface xe2 |
| (config-if) #encapsulation dot1q 1001 | Setting Encapsulation to dot1q with VLAN ID 1001 Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if) #access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if) #map vpn-id 1001 | Map vpn-id 1001 to interface xe2.1001 (VPLS) |
| (config) #commit | Commit candidate configuration to be running configuration |

Validation**PE1: EVPN-VPLS or ELAN**

```

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source           Destination      Status           Up/Down          Update           evpn-id

```

```
=====
10.143.73.1      10.143.73.5      Installed      00:01:03      00:01:03      1001
```

Total number of entries are 1

PE1#sh evpn mpls tunnel label

EVPN-MPLS Network tunnel labels

(*) in Policy - tunnel-policy inherited from mac-vrf

```
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
VP-Multipath      Underlay      Local      Remote      RS
Destination      Status      VPN-ID      Policy      MC-Label      UC-Label      MC-Label      UC-Label      Grp-
Name      NHLFE-ix      NW-Intf      NW-Label
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
10.143.73.5      Installed      1001      --      640      17      640      --      --
--      xe3      24325
```

Total number of entries are 1

PE1#show evpn mpls id 1001

EVPN-MPLS Information

=====

Codes: NW - Network Port

AC - Access Port

(u) - Untagged

```
VPN-ID      EVI-Name      EVI-Type      Type      Interface      ESI      VLAN      DF-Status      Src-
Addr      Dst-Addr
-----
1001      vpls1001      L2      NW      ----      ----      ----      ----
10.143.73.1      10.143.73.5
1001      vpls1001      --      AC      xe6.1001      --- Single Homed Port ---      ----      ----
-----
```

Total number of entries are 2

PE1#sh mpls forwarding-table | inc 10.143.73.5

```
B> 10.143.73.5/32      6      10      -      LSP_
DEFAULT 24960      xe3      No      10.143.73.2
```

PE1#sh ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale

```
Network      Next Hop      In Label      Out Label
*> 10.143.73.1/32 0.0.0.0      24320      -
*>i 10.143.73.2/32 10.143.73.2 24322      24961
*>i 10.143.73.4/32 10.143.73.2 24323      24963
*>i 10.143.73.5/32 10.143.73.2 24321      24960
PE1#
```

PE1#show evpn mpls mac-table

```
=====
=====
EVPN MPLS MAC Entries
=====
=====
VNID      Interface      VlanId      Inner-VlanId      Mac-Addr      VTEP-
Ip/ESI      Type      Status      AccessPortDesc
-----
1001      ----      ----      ----      00ff.2300.0000 10.143.73.5      Dynamic
Remote -----
1001      ----      ----      ----      0023.1001.0001 10.143.73.5      Dynamic
Remote -----
1001      xe6.1001      ----      ----      0045.1001.0001 10.143.73.1      Dynamic
Local -----
```

```

1001      xe6.1001  ----  ----  0071.1001.0001 10.143.73.1      Dynamic
Local    - - - - -  - - - - -

```

PE1: EVPN-VPWS or ELINE

```

PE1#show evpn mpls xconnect
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

```

| Local | | Remote | | Connection-Details | | |
|--------|-------------|--------|-----------|--------------------|---------------------------|-----|
| VPN-ID | EVI-Name | MTU | VPN-ID | Source | Destination | PE- |
| IP | MTU | Type | NW-Status | | | |
| 2 | 10.143.73.5 | 1500 | 252 | xe6.2 | --- Single Homed Port --- | |

Total number of entries are 1

```

PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries

```

| Source | Destination | Status | Up/Down | Update | local-evpn-id | remote- |
|-------------|-------------|-----------|----------|----------|---------------|---------|
| evpn-id | | | | | | |
| 10.143.73.1 | 10.143.73.5 | Installed | 00:01:10 | 00:01:10 | 2 | 252 |

Total number of entries are 1

```

PE1#show evpn mpls xconnect tunnel label
EVPN-MPLS Network tunnel labels
(*) in Policy - tunnel-policy inherited from mac-vrf

```

| Multipath | | Underlay | | Local | Remote | Local | Remote | RSVP- |
|-------------|-----------|----------|----------|--------|----------|----------|--------|-------|
| Destination | Status | VPWS-ID | VPWS-ID | Policy | UC-Label | UC-Label | Grp- | |
| Name | NHLFE-ix | NW-Intf | NW-Label | | | | | |
| 10.143.73.5 | Installed | 2 | 252 | -- | 16 | 16 | -- | -- |
| xe3 | 24325 | | | | | | | |

Total number of entries are 1

```

PE1#sh mpls forwarding-table | inc 10.143.73.5

```

| B> | 10.143.73.5/32 | 6 | 10 | - | - | LSP_ |
|---------|----------------|-----|----|-------------|---|------|
| DEFAULT | 24960 | xe3 | No | 10.143.73.2 | | |

```

PE1#sh ip bgp labeled-unicast

```

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale

| Network | Next Hop | In Label | Out Label |
|--------------------|-------------|----------|-----------|
| *> 10.143.73.1/32 | 0.0.0.0 | 24320 | - |
| *>i 10.143.73.2/32 | 10.143.73.2 | 24322 | 24961 |
| *>i 10.143.73.4/32 | 10.143.73.2 | 24323 | 24963 |
| *>i 10.143.73.5/32 | 10.143.73.2 | 24321 | 24960 |

```
PE1#show evpn mpls xconnect id 2
```

```
EVPN-MPLS Xconnect Info
```

```
=====
```

```
AC-AC: Local-Cross-connect
```

```
AC-NW: Cross-connect to Network
```

```
AC-UP: Access-port is up
```

```
AC-DN: Access-port is down
```

```
NW-UP: Network is up
```

```
NW-DN: Network is down
```

```
NW-SET: Network and AC both are up
```

| Local | | | Remote | | Connection-Details | | |
|--------|-------------|-------|-----------|--------|---------------------------|-----|--|
| VPN-ID | EVI-Name | MTU | VPN-ID | Source | Destination | PE- | |
| IP | MTU | Type | NW-Status | | | | |
| 2 | 10.143.73.5 | 1500 | 252 | xe6.2 | --- Single Homed Port --- | | |
| | | AC-NW | NW-SET | | | | |

```
Total number of entries are 1
```

PE2: EVPN-VPLS or ELAN

```
PE2#show evpn mpls tunnel
```

```
EVPN-MPLS Network tunnel Entries
```

| Source | Destination | Status | Up/Down | Update | evpn-id |
|-------------|-------------|-----------|----------|----------|---------|
| 10.143.73.5 | 10.143.73.1 | Installed | 00:04:03 | 00:04:03 | 1001 |

```
Total number of entries are 1
```

```
PE2#show evpn mpls tunnel label
```

```
EVPN-MPLS Network tunnel labels
```

```
(*) in Policy - tunnel-policy inherited from mac-vrf
```

| VP-Multipath | | | | Underlay | | Local | | Remote | | RS |
|--------------|-----------|---------|----------|----------|----------|----------|----------|--------|--|----|
| Destination | Status | VPN-ID | Policy | MC-Label | UC-Label | MC-Label | UC-Label | Grp- | | |
| Name | NHLFE-ix | NW-Intf | NW-Label | | | | | | | |
| 10.143.73.1 | Installed | 1001 | -- | 640 | 17 | 640 | -- | -- | | |
| -- | xe3 | 24321 | | | | | | | | |

```
Total number of entries are 1
```

```
PE2#
```

```
PE2#sh mpls forwarding-table | include 10.143.73.1/32
```

| B> | 10.143.73.1/32 | 1 | 1 | - | - | LSP_ |
|---------|----------------|-----|----|-------------|---|------|
| DEFAULT | 24961 | xe5 | No | 10.143.73.4 | | |

```
PE2#sh ip bgp labeled-unicast
```

```
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale
```

| Network | Next Hop | In Label | Out Label |
|--------------------|-------------|----------|-----------|
| *>i 10.143.73.1/32 | 10.143.73.4 | 24321 | 24961 |
| *>i 10.143.73.2/32 | 10.143.73.4 | 24323 | 24963 |
| *>i 10.143.73.4/32 | 10.143.73.4 | 24322 | 24960 |
| *> 10.143.73.5/32 | 0.0.0.0 | 24320 | - |

```
PE2#show evpn mpls id 1001
```

```
EVPN-MPLS Information
```

```
=====
```

```
Codes: NW - Network Port
```



```

      AC - Access Port
      (u) - Untagged

VPN-ID   EVI-Name   EVI-Type Type Interface ESI           VLAN DF-Status Src-
Addr           Dst-Addr

1001     vpls1001   L2      NW    ----      ----           ----  ----
      10.143.73.5   10.143.73.1
1001     vpls1001   --      AC    xe2.1001 --- Single Homed Port ---  ----  ----
      ----

```

Total number of entries are 2

PE2#show evpn mpls mac-table

```

=====
EVPN MPLS MAC Entries
=====
VNID      Interface VlanId Inner-VlanId Mac-Addr      VTEP-
Ip/ESI                    Type      Status      AccessPortDesc
-----
1001      ----      ----      ----      0045.1001.0001 10.143.73.1      Dynamic
Remote    -----      -----
1001      ----      ----      ----      0071.1001.0001 10.143.73.1      Dynamic
Remote    -----      -----
1001      xe2.1001 ----      ----      00ff.2300.0000 10.143.73.5      Dynamic
Local     -----      -----
1001      xe2.1001 ----      ----      0023.1001.0001 10.143.73.5      Dynamic
Local     -----      -----

```

PE2: ELINE or EVPN-VPWS

```

PE2#show evpn mpls xconnect
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

```

```

Local      Remote      Connection-Details
=====
VPN-ID      EVI-Name      MTU   VPN-ID      Source      Destination      PE-
IP          MTU   Type   NW-Status
=====
252         ----      1500   2           xe2.2       --- Single Homed Port ---
      10.143.73.1   1500   AC-NW   NW-SET

```

Total number of entries are 1

PE2#show evpn mpls xconnect tunnel

```

EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====
10.143.73.5   10.143.73.1      Installed   00:04:10     00:04:10     252          2

```

Total number of entries are 1

```

PE2#show evpn mpls xconnect tunnel label
EVPN-MPLS Network tunnel labels

```

| Remote | | Local | Remote | Local | | | | |
|-------------|-----------|---------|---------|--------------|--------------|----------|----------|-----|
| Destination | Status | VPWS-ID | VPWS-ID | Network-Intf | Tunnel-Label | MC-Label | UC-Label | MC- |
| Label | UC-Label | | | | | | | |
| 10.143.73.1 | Installed | 252 | 2 | xe5 | 24321 | -- | 16 | - |
| - | 16 | | | | | | | |

```

Total number of entries are 1

PE2#sh mpls forwarding-table | include 10.143.73.1/32
B> 10.143.73.1/32 1 1 - - LSP_
DEFAULT 24961 xe5 No 10.143.73.4

PE2#sh ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale

```

| Network | Next Hop | In Label | Out Label |
|--------------------|-------------|----------|-----------|
| *>i 10.143.73.1/32 | 10.143.73.4 | 24321 | 24961 |
| *>i 10.143.73.2/32 | 10.143.73.4 | 24323 | 24963 |
| *>i 10.143.73.4/32 | 10.143.73.4 | 24322 | 24960 |
| *> 10.143.73.5/32 | 0.0.0.0 | 24320 | - |

```

PE2#show evpn mpls xconnect id 252
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local Remote Connection-Details
=====
=====

```

| VPN-ID | EVI-Name | MTU | VPN-ID | Source | Destination | PE- |
|-------------|----------|-------|-----------|--------|-------------|-----------------------|
| IP | MTU | Type | NW-Status | | | |
| 252 | ---- | 1500 | 2 | xe2.2 | --- | Single Homed Port --- |
| 10.143.73.1 | 1500 | AC-NW | NW-SET | | | |

```

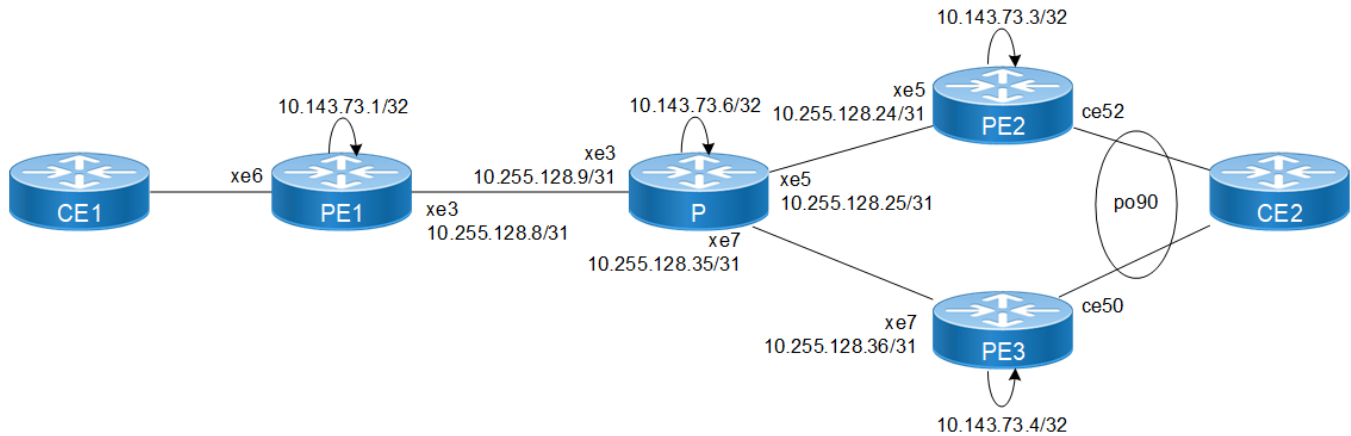
Total number of entries are 1

```

EVPN MPLS Multihoming


Topology

The diagram depicts the Multi Homed topology for the EVPN MPLS configuration examples that follow.

Figure 50. MPLS EVPN multi-homing configuration**PE1: Loopback Interface**

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.1/32 secondary | Configure IP address on loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config-if)#commit | Commit the transaction. |

PE1: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.1 | Set the router ID to IP address 10.143.73.1 |
| (config-router)#transport-address ipv4 10.143.73.1 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| <div>  Note: It is preferable to use the loopback address as the transport address. </div> | |
| (config-router)#targeted-peer ipv4 10.143.73.3 | Configure targeted peer |
| (config-router)#targeted-peer ipv4 10.143.73.4 | Configure targeted peer |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config-router)#commit | Commit the transaction |

PE1: Global EVPN MPLS Command

| | |
|---------------------------|--|
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#commit | Commit candidate configuration to be running configuration |

| | |
|--|---|
| (config)#evpn mpls vtep-ip-global 10.143.73.1 | Configuring VTEP global IP to loopback IP |
| (config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter EVPN-MPLS-MH |
| (config)#evpn mpls multihoming enable | Enable Multihoming |
| (config)#commit | Commit the transaction |

PE1: Interface Configuration Network Side

| | |
|--|---|
| (config)#interface xe3 | Enter the Interface mode for xe3 |
| (config-if)#ip address 10.255.128.8/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config-if)#commit | Commit the transaction |



Note: For RSVP Configuration [RSVP-TE Configuration \(page 643\)](#)

PE1: OSPF Configuration

| | |
|--|--|
| (config)#router ospf 100 | Enter the Router OSPF mode |
| (config-router)#ospf router-id 10.143.73.1 | Router-ID configurations |
| (config-router)#network 10.143.73.1/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.8/31 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit the transaction |

PE1: BGP Configuration

| | |
|---|---|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.3 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.3 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.4 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.4 update-source lo | Source of routing updates as loopback |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.3 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#neighbor 10.143.73.4 activate | Enabling EVPN Address family for neighbor |

| | |
|--------------------------|--------------------------------|
| (config-router-af) #exit | Exiting of Address family mode |
| (config-router) #commit | Commit the transaction |

PE1: MAC VRF Configuration

| | |
|---|--|
| (config) #mac vrf vrf2 | Enter VRF mode |
| (config-vrf) #rd 10.143.73.1:1700 | Configuring Route-Distinguisher value 10.143.73.1:1700 |
| (config-vrf) #route-target both 1700:1700 | Configuring import and export value as 1700:1700 |
| (config-vrf) #exit | Exiting VRF Mode |
| (config) #mac vrf vpls1001 | Enter VRF mode |
| (config-vrf) #rd 10.143.73.1:1001 | Configuring Route-Distinguisher value 10.143.73.1:1001 |
| (config-vrf) #route-target both 1001:1001 | Configuring import and export value as 1001:1001 |
| (config-vrf) #exit | Exiting VRF Mode |
| (config) #commit | Commit the transaction. |

PE1: EVPN-VPWS (ELINE) and VRF Mapping

| | |
|--|---|
| (config) #evpn mpls id 1700 xconnect target-mpls-id 1800 | Configure the EVPN-VPWS identifier with source identifier 1700 and target identifier 1800 |
| (config-evpn-mpls) #host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| (config-evpn-mpls) #commit | Commit the transaction. |
| (config-evpn-mpls) #exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE1:EVPN-VPWS (ELINE) Access Port Configuration

| | |
|--|---|
| (config) #interface xe6 | Enter the Interface mode for xe6. |
| (config-if) #interface xe6.1700 switchport | Creating L2 sub interface of physical interface xe6. |
| (config-if) #description access-side-int | Giving Interface Description |
| (config-if) #encapsulation dot1q 1700 | Setting Encapsulation to dot1q with VLAN ID 1700 Supported. Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if) #access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if) #map vpn-id 1700 | Map vpn-id 1700 to interface xe6.1700 (VPWS). |
| (config-access-if) #exit | Exiting out of access interface mode. |

PE1: EVPN-VPLS (ELAN) and VRF Mapping

| | |
|-----------------------------|--|
| (config) #evpn mpls id 3000 | Configure the EVPN-VPLS identifier with identifier |
|-----------------------------|--|

| | |
|---|---|
| | 3000. |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-VPLS identifier. |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |


PE1: EVPN-VPLS (ELAN) Access Port Configuration

| | |
|--|--|
| (config-if)#interface xe6.300 switchport | Enter the Interface mode for xe6. |
| (config-if)#encapsulation dot1q 3000 | Creating L2 sub interface of physical interface xe6 |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if)#map vpn-id 3000 | Map vpn-id 3000 to interface xe6.3000 (VPLS) |
| (config-access-if)#commit | Commit candidate configuration to be running configuration |

P: Loopback Interface

| | |
|---|---|
| #configure terminal | Enter configure mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.6/32 secondary | Configure IP address on loopback interface |
| (config-if)#exit | Exit interface mode |
| (config-if)#commit | Commit the transaction |

P: Global LDP

| | |
|--|---|
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 10.143.73.6 | Set the router ID to IP address 10.143.73.6 |
| (config-router)#transport-address ipv4 10.143.73.6 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run |
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config-router)#commit | Commit the transaction |

P: Interface Configuration

| | |
|--|---------------------------------------|
| (config)#interface xe3 | Enter the Interface mode for xe3 |
| (config-if)#ip address 10.255.128.9/31 | Configure IP address on the interface |

| | |
|---|---|
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)#ip address 10.255.128.25/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe7 | Enter the Interface mode for po92 |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |

P: OSPF Configuration


| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.6 | Setting the Router ID as Loopback IP |
| (config-router)#network 10.143.73.6/32 area 0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.8/31 area 0.0.0.0 | Advertise xe3 network address in OSPF that comes under same subnet |
| (config-router)#network 10.255.128.24/31 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)#network 10.255.128.36/31 area 0.0.0.0 | Advertise xe7 network address in OSPF that comes under same subnet |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| OcNOS (config)#commit | Commit candidate configuration to be running configuration |

PE2: Loopback Interface

| | |
|---|---|
| #configure terminal | Enter configure mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.3/32 secondary | Configure IP address on loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |

PE2: Global LDP

| | |
|---------------------|----------------------------|
| (config)#router ldp | Enter the Router LDP mode. |
|---------------------|----------------------------|

| | |
|---|---|
| (config-router)#router-id 10.143.73.3 | Set the router ID to IP address 10.143.73.3 |
| (config-router)#transport-address ipv4 10.143.73.3 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run |
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#targeted-peer ipv4 10.143.73.1 | Configure targeted peer |
| (config-router-targeted-peer)# targeted-peer ipv4 10.143.73.4 | Configure targeted peer |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit the transaction |

PE2: Global EVPN MPLS Command

| | |
|--|---|
| (config)#evpn mpls enable | Enable EVPN MPLS Note: Reload is required after Enabling/Disabling EVPN MPLS Feature |
| (config)#commit | Commit candidate configuration to be running configuration |
| (config)#evpn mpls vtep-ip-global 10.143.73.3 | Configuring VTEP global IP to loopback IP |
| (config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter EVPN-MPLS-MH |
| (config)#evpn mpls multihoming enable | Enable Multihoming |
| (config)#evpn esi hold-time 60 | Delay timer for ESI to come up before enabling evpn |
| (config)#commit | Commit the transaction |

PE2: Interface Configuration Network Side

| | |
|---|--|
| (config)#interface xe5 | Enter the Interface mode for xe5. |
| (config-if)#ip address 10.255.128.24/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config-if)#commit | Commit the transaction. |



Note: For RSVP Configuration refer [RSVP-TE Configuration \(page 643\)](#).

PE2: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode |
| (config-router)#ospf router-id 10.143.73.3 | Router-ID configurations |
| (config-router)#network 10.143.73.3/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.24/31 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| OcNOS (config) #commit | Commit candidate configuration to be running configuration |

PE2: BGP Configuration

| | |
|---|---|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.1 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.1 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.4 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.4 update-source lo | Source of routing updates as loopback |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.1 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#neighbor 10.143.73.4 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config-router)#commit | Commit the transaction. |

PE2: MAC VRF Configuration

| | |
|--|--|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.3:1700 | Configuring Route-Distinguisher value 10.143.73.3:1700 |
| (config-vrf)#route-target both 1700:1700 | Configuring import and export value as 1700:1700 Support: route-target export route-target import |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#mac vrf vpls1001 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.3:1001 | Configuring Route-Distinguisher value 10.143.73.3:1001 |
| (config-vrf)#route-target both 1001:1001 | Configuring import and export value as 1001:1001 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#commit | Commit the transaction. |

PE2: EVPN-VPWS (ELINE) and MAC VRF Mapping

| | |
|--|---|
| <code>(config)#evpn mpls id 1800 xconnect target-mpls-id 1700</code> | Configure the EVPN-VPWS identifier with source identifier 1800 and target identifier 1700 |
| <code>(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2</code> | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| <code>(config-evpn-mpls)#commit</code> | Commit the transaction. |
| <code>(config-evpn-mpls)#exit</code> | Exit the EVPN MPLS mode and return to the configure mode. |

PE2:EVPN-VPWS (ELINE) Access Port Configuration

| | |
|---|--|
| <code>(config)#interface po90</code> | Enter the Interface mode for po90. |
| <code>(config-if)#load-interval 30</code> | Load interval setting. |
| <code>(config-if)#evpn multi-homed system-mac 0000.aaaa.bbbc</code> | Configure ESI on a link on which Multi homed CE is connected. |
| <code>(config-if)#interface po90.1700 switchport</code> | Creating L2 sub interface of Dynamic LAG po90. |
| <code>(config-if)#encapsulation dot1q 1700</code> | Setting Encapsulation to dot1q with VLAN ID 1700. Supported Encapsulation: dot1ad, dot1q, untagged, default |
| <code>(config-if)#access-if-evpn</code> | Entering Access mode for EVPN MPLS ID configuration. |
| <code>(config-access-if)#map vpn-id 1800</code> | Map vpn-id 1800 to interface po90.1700 (VPWS). |
| <code>(config-access-if)#exit</code> | Exiting out of access interface mode. |

PE2: EVPN-VPLS (ELAN) and MAC VRF Mapping

| | |
|--|---|
| <code>(config)#evpn mpls id 3000</code> | Configure the EVPN-VPLS identifier with identifier 3000. |
| <code>(config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001</code> | Mapping vrf "vpls1001" to EVPN-VPLS identifier. |
| <code>(config-evpn-mpls)#commit</code> | Commit the transaction. |
| <code>(config-evpn-mpls)#exit</code> | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: EVPN-VPLS (ELAN) Access Port Configuration


| | |
|--|--|
| <code>(config-if)#interface po90.300 switchport</code> | Enter the Interface mode for po90. |
| <code>(config-if)#encapsulation dot1q 3000</code> | Setting Encapsulation to dot1q with VLAN ID 3000. Supported Encapsulation: dot1ad, dot1q, untagged, default |
| <code>(config-if)#access-if-evpn</code> | Entering Access mode for EVPN MPLS ID configuration |

| | |
|--|--|
| (config-access-if)#map vpn-id 3000 | Map vpn-id 3000 to interface po90.300 (VPLS) |
| (config-access-if)#exit | Exiting out of Access if mode. |
| (config-if)#interface ce52 | Enter the Interface mode for ce52. |
| (config-if)#channel-group 90 mode active | Putting interface ce52 in Dynamic LAG po90 |
| (config-access-if)#commit | Commit candidate configuration to be running configuration |

PE3: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.4/32 secondary | Configure IP address on loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |

PE3: Global LDP

| | |
|---|---|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.4 | Set the router ID to IP address 10.143.73.4 |
| (config-router)#transport-address ipv4 10.143.73.4 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#targeted-peer ipv4 10.143.73.1 | Configure targeted peer. |
| (config-router-targeted-peer)# targeted-peer ipv4 10.143.73.3 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit from router LDP mode |
| (config-router)#commit | Commit the transaction. |

PE3: Global EVPN MPLS Command

| | |
|--|--|
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#commit | Commit candidate configuration to be running configuration |
| (config)#evpn mpls vtep-ip-global 10.143.73.4 | Configuring VTEP global IP to loopback IP |
| (config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter EVPN-MPLS-MH |
| (config)#evpn mpls multihoming enable | Enable Multihoming |

| | |
|--------------------------------|---|
| (config)#evpn esi hold-time 60 | Delay timer for ESI to come up before enabling evpn |
| (config)#commit | Commit the transaction. |

PE3: Interface Configuration Network Side

| | |
|---|---|
| (config)#interface xe7 | Enter the Interface mode for xe7. |
| (config-if)#ip address 10.255.128.36/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |



Note: For RSVP Configuration refer [RSVP-TE Configuration \(page 643\)](#)

PE3: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode |
| (config-router)#ospf router-id 10.143.73.4 | Router-ID configurations |
| (config-router)#network 10.143.73.4/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.36/31 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit the transaction |

PE3: BGP Configuration

| | |
|---|---|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.1 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.1 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.3 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.3 update-source lo | Source of routing updates as loopback |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.1 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#neighbor 10.143.73.3 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config-router)#commit | Commit the transaction |

PE3: MAC VRF Configuration

| | |
|--|--|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.4:1700 | Configuring Route-Distinguisher value 10.143.73.4:1700 |
| (config-vrf)#route-target both 1700:1700 | Configuring import and export value as 1700:1700 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#mac vrf vpls1001 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.4:1001 | Configuring Route-Distinguisher value 10.143.73.4:1001 |
| (config-vrf)#route-target both 1001:1001 | Configuring import and export value as 1001:1001 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#commit | Commit the transaction |

PE3: EVPN-VPWS (ELINE) and MAC VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 1800 xconnect target-mpls-id 1700 | Configure the EVPN-VPWS identifier with source identifier 1800 and target identifier 1700 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE3:EVPN-VPWS (ELINE) Access Port Configuration

| | |
|--|---|
| (config)#interface po90 | Enter the Interface mode for po90. |
| (config-if)#load-interval 30 | Load interval setting. |
| (config-if)#evpn multi-homed system-mac 0000.aaaa.bbbc | Configure ESI on a link on which Multi homed CE is connected. |
| (config-if)#interface po90.1700 switchport | Creating L2 sub interface of Dynamic LAG po90. |
| (config-if)#encapsulation dot1q 1700 | Setting Encapsulation to dot1q with VLAN ID 1700 Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if)#map vpn-id 1800 | Map vpn-id 1800 to Dynamic LAG sub interface with po90.1700 (VPWS) |
| (config-access-if)#exit | Exiting out of access interface mode. |

PE3: EVPN-VPLS (ELAN) and MAC VRF Mapping

| | |
|----------------------------|--|
| (config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier with identifier |
|----------------------------|--|

| | |
|---|---|
| | 3000. |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-VPLS identifier. |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE3: EVPN-VPLS (ELAN) Access Port Configuration

| | |
|---|--|
| (config-if)#interface po90.300 switchport | Enter the Interface mode for po90. |
| (config-if)#encapsulation dot1q 3000 | Setting Encapsulation to dot1q with VLAN ID 3000. Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if)#map vpn-id 3000 | Map vpn-id 3000 to interface po90.300 (VPLS) |
| (config-access-if)#exit | Exiting out of Access if mode. |
| (config-if)#interface ce50 | Enter the Interface mode for ce50. |
| (config-if)#channel-group 90 mode active | Putting interface xe2 in Dynamic LAG po90 |
| (config-access-if)#commit | Commit candidate configuration to be running configuration |

Validation

PE1: ELAN or EVPN-VPLS

```

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status          Up/Down         Update          evpn-id
=====
10.143.73.1     10.143.73.4     Installed       00:13:32        00:13:32        3000
10.143.73.1     10.143.73.3     Installed       00:13:33        00:13:33        3000

Total number of entries are 2

PE1#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels

Destination      Status          evpn-id  Network-Intf  Tunnel-Label  Local          Remote
Label           Label           Label    Label         Label         MC-Label  UC-Label  MC-Label  UC-Label
=====
10.143.73.4     Installed       3000     xe3            25601         753         403        654        53
10.143.73.3     Installed       3000     xe3            25600         753         403        753        402

Total number of entries are 2

PE1#show evpn mpls id 3000
EVPN-MPLS Information
=====
Codes: NW - Network Port
      AC - Access Port
      (u) - Untagged

```

| VPN-ID | EVI-Name | EVI-Type | Type | Interface | ESI | VLAN | DF-Status | Src- |
|--------|-------------|-------------|------|-----------|------|-------------------|-----------|------|
| Addr | Dst-Addr | | | | | | | |
| 3000 | ---- | L2 | NW | ---- | ---- | ---- | ---- | ---- |
| | 10.143.73.1 | 10.143.73.4 | | | | | | |
| 3000 | ---- | L2 | NW | ---- | ---- | ---- | ---- | ---- |
| | 10.143.73.1 | 10.143.73.3 | | | | | | |
| 3000 | ---- | -- | AC | xe6.300 | --- | Single Homed Port | --- | ---- |
| | ---- | | | | | | | |

```
PE1#show evpn mpls mac-table
```

```
=====
```

| EVPN MPLS MAC Entries | | | | | |
|-----------------------|-----------|--------|--------------|----------------|-------------------------------|
| ===== | | | | | |
| VNID | Interface | VlanId | Inner-VlanId | Mac-Addr | VTEP- |
| Ip/ESI | | Type | | Status | AccessPortDesc |
| 3000 | xe6.300 | ---- | ---- | 0211.2000.03e8 | 10.143.73.1 |
| Local | ----- | ----- | | | Dynamic |
| 3000 | xe6.300 | ---- | ---- | b86a.97cd.6a3d | 10.143.73.1 |
| Local | ----- | ----- | | | Dynamic |
| 3000 | ---- | ---- | ---- | 0224.2000.03e8 | 00:00:00:aa:aa:bb:bc:00:00:00 |
| Remote | ----- | ----- | | | Dynamic |
| 3000 | ---- | ---- | ---- | b86a.97d2.53bb | 00:00:00:aa:aa:bb:bc:00:00:00 |
| Remote | ----- | ----- | | | Dynamic |

PE1: ELINE or EVPN-VPWS

```
PE1#show evpn mpls xconnect
```

```
EVPN-MPLS Xconnect Info
```

```
=====
```

```
AC-AC: Local-Cross-connect
```

```
AC-NW: Cross-connect to Network
```

```
AC-UP: Access-port is up
```

```
AC-DN: Access-port is down
```

```
NW-UP: Network is up
```

```
NW-DN: Network is down
```

```
NW-SET: Network and AC both are up
```

| Local | | | | Remote | | Connection-Details | |
|-------------|----------|-------|-----------|--------|----------|-------------------------------|-----|
| ===== | | | | ===== | | | |
| | | | | | | | |
| VPN-ID | EVI-Name | | MTU | VPN-ID | Source | Destination | PE- |
| IP | MTU | Type | NW-Status | | | | |
| ===== | | | | ===== | | | |
| | | | | | | | |
| 1700 | ---- | | 1500 | 1800 | xe6.1700 | 00:00:00:aa:aa:bb:bc:00:00:00 | |
| 10.143.73.4 | 1500 | AC-NW | NW-SET | | | | |

```
Total number of entries are 1
```

```
PE1#show evpn mpls xconnect tunnel
```

```
EVPN-MPLS Network tunnel Entries
```

| Source | Destination | Status | Up/Down | Update | local-evpn-id | remote- |
|-------------|-------------|-----------|----------|----------|---------------|---------|
| evpn-id | | | | | | |
| 10.143.73.1 | 10.143.73.4 | Installed | 00:16:50 | 00:16:50 | 1700 | 1800 |
| 10.143.73.1 | 10.143.73.3 | Installed | 00:16:50 | 00:16:50 | 1700 | 1800 |

```
Total number of entries are 2
```

```
PE1#show evpn mpls xconnect tunnel label
```

```
EVPN-MPLS Network tunnel labels
```

| | Local | Remote | Local |
|--------|-------|--------|-------|
| Remote | | | |

```

Destination      Status      VPWS-ID      VPWS-ID      Network-Intf   Tunnel-Label   MC-Label   UC-Label   MC-
Label   UC-Label
=====
10.143.73.4      Installed   1700         1800         xe3             25601          --          402        -
- 52
10.143.73.3      Installed   1700         1800         xe3             25600          --          402        -
- 401

Total number of entries are 2

PE1#show evpn mpls xconnect id 1700
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local              Remote              Connection-Details
=====
VPN-ID      EVI-Name      MTU   VPN-ID      Source      Destination      PE-
IP          MTU   Type   NW-Status
=====
1700      ----      1500  1800      xe6.1700    00:00:00:aa:aa:bb:bc:00:00:00
10.143.73.3      1500  AC-NW  NW-SET
1500      ----      ----
Total number of entries are 1

```

PE2: ELAN or EVPN-VPLS

```

PE2#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      evpn-id
=====
10.143.73.5      10.143.73.6      Installed   00:24:41     00:24:41     3000
10.143.73.5      10.143.73.1      Installed   00:24:38     00:24:38     3000

Total number of entries are 2

PE2#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels

Destination      Status      evpn-id      Network-Intf   Tunnel-Label   Local      Remote
Label            UC-Label
=====
10.143.73.6      Installed   3000         xe5             24965          753        402        654        53
10.143.73.1      Installed   3000         xe5             25604          753        402        753        403

Total number of entries are 2

PE2#show evpn mpls id 3000
EVPN-MPLS Information
=====
Codes: NW - Network Port
       AC - Access Port
       (u) - Untagged

VPN-ID      EVI-Name      EVI-Type Type Interface ESI      VLAN DF-Status Src-
Addr          Dst-Addr
=====

```



```

3000      ----      L2      NW      ----      ----      ----      ----
          10.143.73.5      10.143.73.6
3000      ----      L2      NW      ----      ----      ----      ----
          10.143.73.5      10.143.73.1
3000      ----      --      AC      po90.300  00:00:00:aa:aa:bb:bc:00:00:00  ----  DF      ----
          ----

```

Total number of entries are 3

PE2#show evpn mpls mac-table

```

=====
EVPN MPLS MAC Entries
=====
VNID      Interface VlanId Inner-VlanId Mac-Addr      VTEP-
Ip/ESI                                Status      AccessPortDesc
-----
3000      ----      ----      ----      0211.2000.03e8 10.143.73.1      Dynamic
Remote   -----      -----
3000      ----      ----      ----      b86a.97cd.6a3d 10.143.73.1      Dynamic
Remote   -----      -----
3000      ----      ----      ----      0224.2000.03e8 00:00:00:aa:aa:bb:bc:00:00:00  Dynamic
Remote   -----      -----
3000      po90.300  ----      ----      b86a.97d2.53bb 00:00:00:aa:aa:bb:bc:00:00:00  Dynamic
Local    -----      -----

```

PE2: ELINE or EVPN-VPWS

PE2#show evpn mpls xconnect id 1800

EVPN-MPLS Xconnect Info

=====

AC-AC: Local-Cross-connect

AC-NW: Cross-connect to Network

AC-UP: Access-port is up

AC-DN: Access-port is down

NW-UP: Network is up

NW-DN: Network is down

NW-SET: Network and AC both are up

```

Local      Remote      Connection-Details
=====
VPN-ID      EVI-Name      MTU      VPN-ID      Source      Destination      PE-
IP          MTU      Type      NW-Status
=====
1800      ----      1500      1700      po90.1700      --- Single Homed Port ---
          10.143.73.1      1500      AC-NW      NW-SET

```

Total number of entries are 1

PE2#show evpn mpls xconnect tunnel

EVPN-MPLS Network tunnel Entries

```

Source      Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====

```

```

10.143.73.3      10.143.73.1      Installed      00:50:18      00:50:18      1800      1700

```

Total number of entries are 1

PE2#show evpn mpls xconnect

EVPN-MPLS Xconnect Info

=====

AC-AC: Local-Cross-connect

AC-NW: Cross-connect to Network

```

AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

```

```

Local                               Remote      Connection-Details
=====
VPN-ID      EVI-Name      MTU  VPN-ID      Source      Destination      PE-
IP          MTU    Type    NW-Status
=====
1800        ----        1500  1700        po90.1700    --- Single Homed Port ---
      10.143.73.1      1500  AC-NW  NW-SET

```

Total number of entries are 1

```

PE2#show evpn mpls xconnect tunnel label
EVPN-MPLS Network tunnel labels

```

```

Remote                               Local      Remote                               Local
Destination      Status      VPWS-ID  VPWS-ID  Network-Intf  Tunnel-Label  MC-Label  UC-Label  MC-
Label  UC-Label
=====
10.143.73.1      Installed  1800     1700     xe5            25604         --         401      -
-         402

```

Total number of entries are 1

PE3: ELAN or EVPN-VPLS

```

PE3#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries

```

```

Source      Destination      Status      Up/Down      Update      evpn-id
=====
10.143.73.4  10.143.73.3      Installed    00:22:11     00:22:11     3000
10.143.73.4  10.143.73.1      Installed    00:22:11     00:22:11     3000

```

Total number of entries are 2

```

PE3#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels

```

```

Destination      Status      evpn-id  Network-Intf  Tunnel-Label  Local      Remote
Label            UC-Label
=====
10.143.73.3      Installed  3000     xe7            24962         654        53        753        402
10.143.73.1      Installed  3000     xe7            24964         654        53        753        403

```

Total number of entries are 2

```

PE3#show evpn mpls id 3000

```

```

EVPN-MPLS Information

```

```

=====

```

```

Codes: NW - Network Port
       AC - Access Port
       (u) - Untagged

```

```

VPN-ID      EVI-Name      EVI-Type  Type  Interface  ESI      VLAN  DF-Status  Src-
Addr          Dst-Addr
=====
3000        ----        L2        NW    ----        ----        ----  ----
      10.143.73.4      10.143.73.3

```

```

3000      ----      L2      NW      ----      ----      ----      ----
          10.143.73.4      10.143.73.1
3000      ----      --      AC      po90.300      00:00:00:aa:aa:bb:bc:00:00:00      ----      NON-DF      ----
          ----

```

Total number of entries are 3

PE3#show evpn mpls mac-table

```

=====
=====
                                EVPN MPLS MAC Entries
=====
=====
VNID      Interface VlanId Inner-VlanId Mac-Addr      VTEP-
Ip/ESI                                Status      AccessPortDesc
-----
3000      ----      ----      ----      0211.2000.03e8 10.143.73.1      Dynamic
Remote
3000      ----      ----      ----      b86a.97cd.6a3d 10.143.73.1      Dynamic
Remote
3000      ----      ----      ----      0224.2000.03e8 00:00:00:aa:aa:bb:bc:00:00:00      Dynamic
Remote
3000      po90.300      ----      ----      b86a.97d2.53bb 00:00:00:aa:aa:bb:bc:00:00:00      Dynamic
Local

```

PE3: ELINE or EVPN-VPWS

PE3#show evpn mpls xconnect id 1800

EVPN-MPLS Xconnect Info

=====

AC-AC: Local-Cross-connect

AC-NW: Cross-connect to Network

AC-UP: Access-port is up

AC-DN: Access-port is down

NW-UP: Network is up

NW-DN: Network is down

NW-SET: Network and AC both are up

```

Local      Remote      Connection-Details
=====
=====
VPN-ID      EVI-Name      MTU      VPN-ID      Source      Destination      PE-
IP      MTU      Type      NW-Status
=====
1800      ----      1500      1700      po90.1700      --- Single Homed Port ---
          10.143.73.1      1500      AC-NW      NW-SET

```

Total number of entries are 1

PE3#show evpn mpls xconnect tunnel

EVPN-MPLS Network tunnel Entries

```

Source      Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====

```

```

10.143.73.4      10.143.73.1      Installed      00:23:18      00:23:18      1800      1700

```

Total number of entries are 1

PE3#show evpn mpls xconnect tunnel label

EVPN-MPLS Network tunnel labels

```

Remote      Local      Remote      Local
Destination      Status      VPWS-ID      VPWS-ID      Network-Intf      Tunnel-Label      MC-Label      UC-Label      MC-
Label      UC-Label
=====

```

```
=====
10.143.73.1      Installed    1800    1700      xe7      24964      --
   52           --          402

Total number of entries are 1

PE3#show evpn mpls xconnect
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                      Remote      Connection-Details
=====
=====
VPN-ID      EVI-Name      MTU  VPN-ID      Source      Destination      PE-
IP          MTU   Type  NW-Status
=====
=====
1800        -----      1500  1700        po90.1700    --- Single Homed Port ---
   10.143.73.1      1500  AC-NW  NW-SET

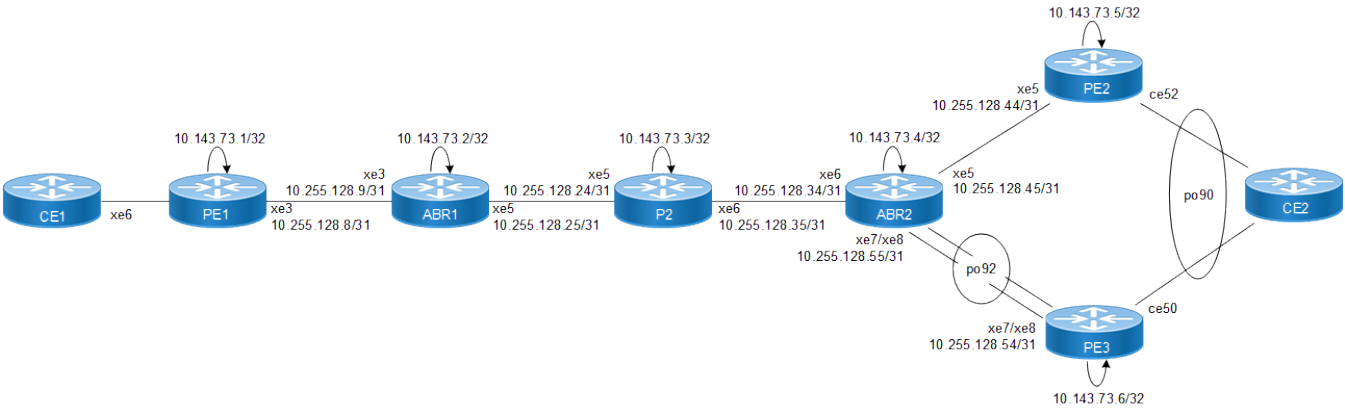
Total number of entries are 1
```

EVPN MPLS Multihoming over BGP-LU

Topology

The diagram depicts the Multi Homed topology for the EVPN MPLS configuration and examples for both ELINE and ELAN service with BGP-LU as underlay MPLS path.

Figure 51. MPLS EVPN multi-homing over LU configuration




Configuration

PE1: Loopback Interface

| | |
|-----------------------|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |

| | |
|---|---|
| (config-if)#ip address 10.143.73.1/32 secondary | Configure IP address on loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |

PE1: Global LDP

| | |
|---|---|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.1 | Set the router ID to IP address 10.143.73.1 |
| (config-router)#transport-address ipv4 10.143.73.1 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run.  Note: It is preferable to use the loopback address as the transport address. |
| (config-router-targeted-peer)# targeted-peer ipv4 10.143.73.6 | Configure targeted peer |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit the transaction |

PE1: Global EVPN MPLS Command

| | |
|--|--|
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#commit | Commit candidate configuration to be running configuration |
| (config)#evpn mpls vtep-ip-global 10.143.73.1 | Configuring VTEP global IP to loopback IP |
| (config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter EVPN-MPLS-MH |
| (config)#evpn mpls multihoming enable | Enable Multihoming |
| (config)#commit | Commit the transaction |

PE1: Interface Configuration Network Side

| | |
|--|---|
| (config)#interface xe3 | Enter the Interface mode for xe3 |
| (config-if)#ip address 10.255.128.8/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |

**Note:** For RSVP Configuration refer [Configure RSVP-TE \(page 644\)](#).

PE1: OSPF Configuration

| | |
|--|--|
| (config)#router ospf 100 | Enter the Router OSPF mode |
| (config-router)#ospf router-id 10.143.73.1 | Router-ID configurations |
| (config-router)#network 10.143.73.1/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.8/31 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit the transaction |

PE1: BGP Configuration

| | |
|---|---|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.5 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.5 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.6 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.6 update-source lo | Source of routing updates as loopback |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.5 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#neighbor 10.143.73.6 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config-router)#commit | Commit the transaction |

PE1: MAC VRF Configuration

| | |
|--|--|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.1:1700 | Configuring Route-Distinguisher value 10.143.73.1:1700 |
| (config-vrf)#route-target both 1700:1700 | Configuring import and export value as 1700:1700 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#mac vrf vpls1001 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.1:1001 | Configuring Route-Distinguisher value 10.143.73.1:1001 |
| (config-vrf)#route-target both 1001:1001 | Configuring import and export value as 1001:1001 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#commit | Commit the transaction |

PE1: EVPN-VPWS (ELINE) and VRF Mapping

| | |
|--|---|
| <code>(config)#evpn mpls id 1700 xconnect target-mpls-id 1800</code> | Configure the EVPN-VPWS identifier with source identifier 1700 and target identifier 1800 |
| <code>(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2</code> | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| <code>(config-evpn-mpls)#commit</code> | Commit the transaction. |
| <code>(config-evpn-mpls)#exit</code> | Exit the EVPN MPLS mode and return to the configure mode. |

PE1:EVPN-VPWS (ELINE) Access Port Configuration

| | |
|--|---|
| <code>(config)#interface xe6</code> | Enter the Interface mode for xe6. |
| <code>(config-if)#interface xe6.1700 switchport</code> | Creating L2 sub interface of physical interface xe6. |
| <code>(config-if)#description access-side-int</code> | Giving Interface Description |
| <code>(config-if)#encapsulation dot1q 1700</code> | Setting Encapsulation to dot1q with VLAN ID 1700 Supported. Encapsulation: dot1ad, dot1q, untagged, default |
| <code>(config-if)#access-if-evpn</code> | Entering Access mode for EVPN MPLS ID configuration. |
| <code>(config-access-if)#map vpn-id 1700</code> | Map vpn-id 1700 to interface xe6.1700 (VPWS). |
| <code>(config-access-if)#exit</code> | Exiting out of access interface mode. |

PE1: EVPN-VPLS (ELAN) and VRF Mapping

| | |
|--|---|
| <code>(config)#evpn mpls id 3000</code> | Configure the EVPN-VPLS identifier with identifier 3000. |
| <code>(config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001</code> | Mapping vrf "vpls1001" to EVPN-VPLS identifier. |
| <code>(config-evpn-mpls)#commit</code> | Commit the transaction. |
| <code>(config-evpn-mpls)#exit</code> | Exit the EVPN MPLS mode and return to the configure mode. |

PE1: EVPN-VPLS (ELAN) Access Port Configuration

| | |
|---|--|
| <code>(config-if)#interface xe6.300 switchport</code> | Enter the Interface mode for xe6. |
| <code>(config-if)#encapsulation dot1q 3000</code> | Creating L2 sub interface of physical interface xe6 |
| <code>(config-if)#access-if-evpn</code> | Entering Access mode for EVPN MPLS ID configuration |
| <code>(config-access-if)#map vpn-id 3000</code> | Map vpn-id 3000 to interface xe6.3000 (VPLS) |
| <code>(config-access-if)#commit</code> | Commit candidate configuration to be running configuration |

ABR1: Loopback Interface

| | |
|---|---|
| #configure terminal | Enter configuration mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.2/32 secondary | Configure IP address on loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |

ABR1: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 10.143.73.2 | Set the router ID to IP address 10.143.73.2 |
| (config-router)#transport-address ipv4 10.143.73.2 | Configure transport address under router ldp |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit the transaction |

ABR1: Interface Configuration

| | |
|---|---|
| (config)#interface xe3 | Enter the Interface mode for xe3 |
| (config-if)#ip address 10.255.128.9/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)#ip address 10.255.128.25/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |

ABR1: OSPF Configuration

| | |
|---|--------------------------------------|
| (config)#enable ext-ospf-multi-inst | Enable multiple-instance capabilit |
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.2 | Setting the Router ID as Loopback IP |
| (config-router)#network 10.143.73.2/32 area 0.0.0.0 instance-id 200 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.25/31 area 0.0.0.0 | Advertise network address in OSPF |
| (config)#router ospf 200 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.2 | Setting the Router ID as Loopback IP |

| | |
|---|--|
| (config-router)#network 10.143.73.2/32 area 0.0.0.0 instance-id 100 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.9/31 area 0.0.0.0 | Advertise network address in OSPF that comes under same subnet |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit candidate configuration to be running configuration |

ABR1: BGP-LU Configuration

| | |
|---|--|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config)#allocate-label all | Configure allocate all under router BGP |
| (config-router)#neighbor 10.143.73.1 remote-as 65010 | Configuring PE1 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.4 remote-as 65010 | Configuring ABR2 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.1 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.4 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 unicast | Entering into address family ipv4 unicast |
| (config-router-af)# network 10.143.73.2/32 | Advertise the network |
| (config-router-af)# exit-address-family | Exit from IPv4 address family |
| (config-router)# address-family ipv4 labeled-unicast | Entering into address family ipv4 labeled-unicast |
| (config-router-af)#neighbor 10.143.73.4 activate | Activate the ipv4 labeled-unicast neighbor |
| (config-router-af)#neighbor 10.143.73.4 next-hop-self | Configure next-hopself for the ipv4 labeled-unicast neighbor |
| config-router-af)#neighbor 10.143.73.4 route-reflector-client | Configure neighbor as route reflector client |
| (config-router-af)#neighbor 10.143.73.1 activate | Activate the ipv4 labeled-unicast neighbor |
| config-router-af)#neighbor 10.143.73.1 route-reflector-client | Configure neighbor as route reflector client |
| (config-router-af)#neighbor 10.143.73.1 next-hop-self | Configure next-hopself for the ipv4 labeled-unicast neighbor |
| (config-router-af)#commit | Commit the transaction |

P2: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 10.143.73.3/32 secondary | Configure IP address on loopback interface. |

| | |
|------------------|------------------------|
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |

P2: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.3 | Set the router ID to IP address 10.143.73.3 |
| (config-router)#transport-address ipv4 10.143.73.3 | Configure transport address under router LDP |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

P2: Interface Configuration

| | |
|---|--|
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)#ip address 10.255.128.24/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe6 | Enter the Interface mode for xe6 |
| (config-if)#ip address 10.255.128.35/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |

P2: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 200 | Enter the Router OSPF mode |
| (config-router)#ospf router-id 10.143.73.3 | Setting the Router ID as Loopback IP |
| (config-router)#network 10.143.73.3/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.24/31 area 0.0.0.0 | Advertise network address in OSPF that comes under same subnet |
| (config-router)#network 10.255.128.35/31 area 0.0.0.0 | Advertise xe5 network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit candidate configuration to be running configuration |

ABR2: Loopback Interface

| | |
|---|---|
| #configure terminal | Enter configuration mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.4/32 secondary | Configure IP address on loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |

ABR2: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 10.143.73.4 | Set the router ID to IP address 10.143.73.4 |
| (config-router)#transport-address ipv4 10.143.73.4 | Configure transport address under router ldp |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit the transaction |

ABR2: Interface Configuration

| | |
|--|---|
| (config)#interface xe6 | Enter the Interface mode for xe3 |
| (config-if)#ip address 10.255.128.34/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)#ip address 10.255.128.45/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#interface po92 | Enter the Interface mode for po92 |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#ip address 10.255.128.55/31 | Configure IP address on the interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#interface xe7 | Enter the Interface mode for xe7 |
| (config-if)#channel-group 92 mode active | Moving interface to Dynamic LAG 92 |
| (config-if)#interface xe8 | Enter the Interface mode for xe8 |
| (config-if)#channel-group 92 mode active | Moving interface to Dynamic LAG 92 |
| (config-if)#commit | Commit the transaction |

ABR2: OSPF Configuration

| | |
|---|---|
| (config)#enable ext-ospf-multi-inst | Enable multiple-instance capabilit |
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.4 | Setting the Router ID as Loopback IP |
| (config-router)#network 10.143.73.4/32 area 0.0.0.0 instance-id 300 | Advertise loopback address in OSPF. |
| (config-router)#network 10.255.128.34/31 area 0.0.0.0 | Advertise network address in OSPF that comes under same subnet. |
| (config)#router ospf 300 | Enter the Router OSPF mode. |
| (config-router)#network 10.143.73.4/32 area 0.0.0.0 instance-id 200 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.45/31 area 0.0.0.0 | Advertise xe5 network address in OSPF |
| (config-router)#network 10.255.128.55/31 area 0.0.0.0 | Advertise po network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit candidate configuration to be running configuration |

ABR2: BGP-LU Configuration

| | |
|---|--|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config)#allocate-label all | Configure allocate all under router BGP |
| (config-router)#neighbor 10.143.73.2 remote-as 65010 | Configuring ABR1 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.5 remote-as 65010 | Configuring PE2 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.6 remote-as 65010 | Configuring PE3 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.2 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.5 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.6 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 unicast | Entering into address family ipv4 unicast |
| (config-router-af)# network 10.143.73.4/32 | Advertise the network |
| (config-router-af)# exit-address-family | Exit from IPv4 address family |
| (config-router)# address-family ipv4 labeled-unicast | Entering into address family ipv4 labeled-unicast |
| (config-router-af)#neighbor 10.143.73.2 activate | Activate the ipv4 labeled-unicast neighbor |
| (config-router-af)#neighbor 10.143.73.2 next-hop- | Configure next-hop-self for the ipv4 labeled-unicast |

| | |
|---|---|
| self | neighbor |
| config-router-af)#neighbor 10.143.73.2 route-reflector-client | |
| (config-router-af)#neighbor 10.143.73.5 activate | Activate the ipv4 labeled-unicast neighbor |
| config-router-af)#neighbor 10.143.73.5 route-reflector-client | Configure neighbor as route reflector client |
| (config-router-af)#neighbor 10.143.73.5 next-hop-self | Configure next-hop-self for the ipv4 labeled-unicast neighbor |
| (config-router-af)#neighbor 10.143.73.6 activate | Activate the ipv4 labeled-unicast neighbor |
| config-router-af)#neighbor 10.143.73.6 route-reflector-client | Configure neighbor as route reflector client |
| (config-router-af)#neighbor 10.143.73.6 next-hop-self | Configure next-hop-self for the ipv4 labeled-unicast neighbor |
| (config-router-af)#commit | Commit candidate configuration to be running configuration |

PE2: Loopback Interface


| | |
|---|--|
| #configure terminal | Enter configure mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.5/32 secondary | Configure IP address on loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 10.143.73.5 | Set the router ID to IP address 10.143.73.5 |
| (config-router)#transport-address ipv4 10.143.73.5 | Configure transport address under router ldp |
| (config-router)#targeted-peer ipv4 10.143.73.1 | Configure targeted peer |
| (config-router-targeted-peer)#targeted-peer ipv4 10.143.73.6 | Configure targeted peer |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: Global EVPN MPLS Command

| | |
|---------------------------|------------------|
| (config)#evpn mpls enable | Enable EVPN MPLS |
|---------------------------|------------------|

| | |
|--|--|
| |  Note: Reload is required after Enabling or Disabling EVPN MPLS. |
| (config)#commit | Commit candidate configuration to be running configuration |
| (config)#evpn mpls vtep-ip-global 10.143.73.5 | Configuring VTEP global IP to loopback IP |
| (config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter EVPN-MPLS-MH |
| (config)#evpn mpls multihoming enable | Enable Multihoming |
| (config)#evpn esi hold-time 60 | Delay timer for ESI to come up before enabling evpn |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: Interface Configuration Network Side

| | |
|---|--|
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)#ip address 10.255.128.44/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |



Note: For RSVP Configuration refer [RSVP-TE Configuration \(page 643\)](#).

PE2: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 300 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.5 | Router-ID configurations |
| (config-router)#network 10.143.73.5/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.44/31 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: BGP Configuration

| | |
|--|---------------------------------------|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.1 remote-as 65010 | Configuring Neighbor as iBGP neighbor |

| | |
|---|---|
| (config-router)#neighbor 10.143.73.1 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.4 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.4 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.6 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.6 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 labeled-unicast | Entering into address family ipv4 labeled-unicast |
| (config-router-af)#neighbor 10.143.73.4 activate | Activate the ipv4 labeled-unicast neighbor |
| (config-router-af)#neighbor 10.143.73.4 next-hop-self | Configure next-hop-self for the ipv4 labeled-unicast neighbor |
| (config-router-af)#exit-address-family | Exit the address family |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.1 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#neighbor 10.143.73.6 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config-router)#commit | Commit candidate configuration to be running configuration |

PE2: MAC VRF Configuration

| | |
|--|--|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.5:1700 | Configuring Route-Distinguisher value 10.143.73.5:1700 |
| (config-vrf)#route-target both 1700:1700 | Configuring import and export value as 1700:1700 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#mac vrf vpls1001 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.5:1001 | Configuring Route-Distinguisher value 10.143.73.5:1001 |
| (config-vrf)#route-target both 1001:1001 | Configuring import and export value as 1001:1001 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: EVPN-VPWS (ELINE) and MAC VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 1800 xconnect target-mpls-id 1700 | Configure the EVPN-VPWS identifier with source identifier 1800 and target identifier 1700 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier |

| | |
|----------------------------|---|
| (config-evpn-mpls) #commit | Commit the transaction. |
| (config-evpn-mpls) #exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: EVPN-VPWS (ELINE) Access Port Configuration

| | |
|---|--|
| (config) #interface po90 | Enter the Interface mode for po90. |
| (config-if) #load-interval 30 | Load interval setting. |
| (config-if) #evpn multi-homed system-mac 0000.aaaa.bbbc | Configure ESI on a link on which Multi homed CE is connected. |
| (config-if) #interface po90.1700 switchport | Creating L2 sub interface of Dynamic LAG po90. |
| (config-if) #encapsulation dot1q 1700 | Setting Encapsulation to dot1q with VLAN ID 1700. Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if) #access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| (config-access-if) #map vpn-id 1800 | Map vpn-id 1800 to interface po90.1700 (VPWS). |
| (config-access-if) #exit | Exiting out of access interface mode. |

PE2: EVPN-VPLS (ELAN) and MAC VRF Mapping

| | |
|--|---|
| (config) #evpn mpls id 3000 | Configure the EVPN-VPLS identifier with identifier 3000. |
| (config-evpn-mpls) #host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-VPLS identifier. |
| (config-evpn-mpls) #commit | Commit the transaction. |
| (config-evpn-mpls) #exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: EVPN-VPLS (ELAN) Access Port Configuration

| | |
|--|--|
| (config-if) #interface po90.300 switchport | Enter the Interface mode for po90. |
| (config-if) #encapsulation dot1q 3000 | Setting Encapsulation to dot1q with VLAN ID 3000. Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if) #access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if) #map vpn-id 3000 | Map vpn-id 3000 to interface po90.300 (VPLS) |
| (config-access-if) #exit | Exiting out of Access if mode. |
| (config-if) #interface ce52 | Enter the Interface mode for ce52. |
| (config-if) #channel-group 90 mode active | Putting interface ce52 in Dynamic LAG po90 |
| (config-access-if) #commit | Commit candidate configuration to be running configuration |

PE3: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configure mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.6/32 secondary | Configure IP address on loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE3: Global LDP

| | |
|--|--|
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 10.143.73.6 | Set the router ID to IP address 10.143.73.5 |
| (config-router)#transport-address ipv4 10.143.73.6 | Configure transport address under router ldp |
| (config-router)#targeted-peer ipv4 10.143.73.1 | Configure targeted peer |
| (config-router-targeted-peer)#targeted-peer ipv4 10.143.73.5 | Configure targeted peer |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit from router LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE3: Global EVPN MPLS Command

| | |
|--|--|
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#commit | Commit candidate configuration to be running configuration |
| (config)#evpn mpls vtep-ip-global 10.143.73.6 | Configuring VTEP global IP to loopback IP |
| (config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter EVPN-MPLS-MH |
| (config)#evpn mpls multihoming enable | Enable Multihoming |
| (config)#evpn esi hold-time 60 | Delay timer for ESI to come up before enabling evpn |
| (config)#commit | Commit candidate configuration to be running configuration |

PE3: Interface Configuration Network Side

| | |
|---|---|
| (config)#interface po92 | Enter the Interface mode for xe3 |
| (config-if)#ip address 10.255.128.54/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |

| | |
|--|--|
| (config-if)#interface xe7 | Enter the Interface mode for xe7 |
| (config-if)#channel-group 92 mode active | Moving interface to Dynamic LAG 92 |
| (config-if)#interface xe8 | Enter the Interface mode for xe8 |
| (config-if)#channel-group 92 mode active | Moving interface to Dynamic LAG 92 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE3: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 300 | Enter the Router OSPF mode |
| (config-router)#ospf router-id 10.143.73.6 | Router-ID configurations |
| (config-router)#network 10.143.73.6/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.54/31 area 0.0.0.0 | Advertise network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE3: BGP Configuration

| | |
|---|---|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.1 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.1 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.4 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.4 update-source lo | Source of routing updates as loopback |
| (config-router)#neighbor 10.143.73.5 remote-as 65010 | Configuring Neighbor as iBGP neighbor |
| (config-router)#neighbor 10.143.73.5 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 labeled-unicast | Entering into address family ipv4 labeled-unicast |
| (config-router-af)#neighbor 10.143.73.4 activate | Activate the ipv4 labeled-unicast neighbor |
| (config-router-af)#neighbor 10.143.73.4 next-hop-self | Configure next-hop-self for the ipv4 labeled-unicast neighbor |
| (config-router-af)#exit-address-family | Exit the address family |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.1 activate | Enabling EVPN Address family for neighbor |

| | |
|--|--|
| (config-router-af)#neighbor 10.143.73.5 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config-router)#commit | Commit candidate configuration to be running configuration |

PE3: MAC VRF Configuration

| | |
|--|--|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.6:1700 | Configuring Route-Distinguisher value 10.143.73.6:1700 |
| (config-vrf)#route-target both 1700:1700 | Configuring import and export value as 1700:1700 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#mac vrf vpls1001 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.6:1001 | Configuring Route-Distinguisher value 10.143.73.6:1001 |
| (config-vrf)#route-target both 1001:1001 | Configuring import and export value as 1001:1001 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE3: EVPN-VPWS (ELINE) and VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 1800 xconnect target-mpls-id 1700 | Configure the EVPN-VPWS identifier with source identifier 1800 and target identifier 1700 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE3: EVPN-VPWS (ELINE) Access Port Configuration

| | |
|--|---|
| (config)#interface po90 | Enter the Interface mode for po90. |
| (config-if)#load-interval 30 | Load interval setting. |
| (config-if)#evpn multi-homed system-mac 0000.aaaa.bbbc | Configure ESI on a link on which Multi homed CE is connected. |
| (config-if)#interface po90.1700 switchport | Creating L2 sub interface of Dynamic LAG po90. |
| (config-if)#encapsulation dot1q 1700 | Setting Encapsulation to dot1q with VLAN ID 1700 Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |

| | |
|------------------------------------|--|
| (config-access-if)#map vpn-id 1800 | Map vpn-id 1800 to Dynamic LAG sub interface with po90.1700 (VPWS) |
| (config-access-if)#exit | Exiting out of access interface mode. |

PE3: EVPN-VPLS (ELAN) and VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier with identifier 3000. |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-VPLS identifier. |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE3: EVPN-VPLS (ELAN) Access Port Configuration

| | |
|---|--|
| (config-if)#interface po90.300 switchport | Enter the Interface mode for po90. |
| (config-if)#encapsulation dot1q 3000 | Setting Encapsulation to dot1q with VLAN ID 3000. Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if)#map vpn-id 3000 | Map vpn-id 3000 to interface po90.300 (VPLS) |
| (config-access-if)#exit | Exiting out of Access if mode. |
| (config-if)#interface ce50 | Enter the Interface mode for ce50. |
| (config-if)#channel-group 90 mode active | Putting interface ce50 in Dynamic LAG po90 |
| (config-access-if)#commit | Commit candidate configuration to be running configuration |

Validation

PE1: ELAN or EVPN-VPLS

```

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status          Up/Down          Update          evpn-id
=====
10.143.73.1     10.143.73.5     Installed       00:13:32         00:13:32         3000
10.143.73.1     10.143.73.6     Installed       00:13:33         00:13:33         3000

Total number of entries are 2

PE1#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels
Destination      Status          evpn-id  Network-Intf  Tunnel-Label  Local          Remote
Label
=====
10.143.73.5     Installed       3000     xe3           25601         753           403           654           53

```

```
10.143.73.6      Installed   3000      xe3        25600      753        403        753        402
```

Total number of entries are 2

```
PE1#sh mpls forwarding-table | inc 10.143.73.5
```

```
B> 10.143.73.5/32      6      10      -      -      LSP_
DEFAULT 25601          xe3          No    10.143.73.2
```

```
PE1#sh mpls forwarding-table | inc 10.143.73.6
```

```
B> 10.143.73.6/32      6      10      -      -      LSP_
DEFAULT 25660          xe3          No    10.143.73.2
```

```
PE1#sh ip bgp labeled-unicast
```

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, S - stale

| Network | Next Hop | In Label | Out Label |
|--------------------|-------------|----------|-----------|
| *> 10.143.73.1/32 | 0.0.0.0 | 24320 | - |
| *>i 10.143.73.2/32 | 10.143.73.2 | 24322 | 25603 |
| *>i 10.143.73.4/32 | 10.143.73.2 | 24323 | 25602 |
| *>i 10.143.73.5/32 | 10.143.73.2 | 24321 | 25601 |
| *>i 10.143.73.6/32 | 10.143.73.2 | 24324 | 25600 |

```
PE1#show evpn mpls id 3000
```

EVPN-MPLS Information

=====

Codes: NW - Network Port
AC - Access Port
(u) - Untagged

| VPN-ID | EVI-Name | EVI-Type | Type | Interface | ESI | VLAN | DF-Status | Src- |
|--------|-------------|----------|-------------|-----------|------|-------------------|-----------|------|
| Addr | Dst-Addr | | | | | | | |
| 3000 | ---- | L2 | NW | ---- | ---- | | | |
| | 10.143.73.1 | | 10.143.73.5 | | | | | |
| 3000 | ---- | L2 | NW | ---- | ---- | | | |
| | 10.143.73.1 | | 10.143.73.6 | | | | | |
| 3000 | ---- | -- | AC | xe6.300 | --- | Single Homed Port | --- | ---- |
| | ---- | | | | | | | |

```
PE1#show evpn mpls mac-table
```

=====

EVPN MPLS MAC Entries

=====

| VNID | Interface | VlanId | Inner-VlanId | Mac-Addr | VTEP- |
|--------|-----------|--------|--------------|----------------|-------------------------------|
| Ip/ESI | | Type | | Status | AccessPortDesc |
| 3000 | xe6.300 | ---- | ---- | 0211.2000.03e8 | 10.143.73.1 |
| Local | ----- | ----- | | | Dynamic |
| 3000 | xe6.300 | ---- | ---- | b86a.97cd.6a3d | 10.143.73.1 |
| Local | ----- | ----- | | | Dynamic |
| 3000 | ---- | ---- | ---- | 0224.2000.03e8 | 00:00:00:aa:aa:bb:bc:00:00:00 |
| Remote | ----- | ----- | | | Dynamic |
| 3000 | ---- | ---- | ---- | b86a.97d2.53bb | 00:00:00:aa:aa:bb:bc:00:00:00 |
| Remote | ----- | ----- | | | Dynamic |

| Network | Next Hop | In Label | Out Label |
|--------------------|-------------|----------|-----------|
| *> 10.143.73.1/32 | 0.0.0.0 | 24320 | - |
| *>i 10.143.73.2/32 | 10.143.73.2 | 24322 | 25603 |
| *>i 10.143.73.4/32 | 10.143.73.2 | 24323 | 25602 |

PE1: ELINE or EVPN-VPWS

```
PE1#show evpn mpls xconnect
```

```
EVPN-MPLS Xconnect Info
```

```
=====
```

```
AC-AC: Local-Cross-connect
```

```
AC-NW: Cross-connect to Network
```

```
AC-UP: Access-port is up
```

```
AC-DN: Access-port is down
```

```
NW-UP: Network is up
```

```
NW-DN: Network is down
```

```
NW-SET: Network and AC both are up
```

| Local | | | Remote | | Connection-Details | | |
|-------------|----------|-------|-----------|----------|-------------------------------|-----|--|
| VPN-ID | EVI-Name | MTU | VPN-ID | Source | Destination | PE- | |
| IP | MTU | Type | NW-Status | | | | |
| 1700 | ---- | 1500 | 1800 | xe6.1700 | 00:00:00:aa:aa:bb:bc:00:00:00 | | |
| 10.143.73.5 | 1500 | AC-NW | NW-SET | | | | |

```
Total number of entries are 1
```

```
PE1#show evpn mpls xconnect tunnel
```

```
EVPN-MPLS Network tunnel Entries
```

| Source | Destination | Status | Up/Down | Update | local-evpn-id | remote-evpn-id |
|-------------|-------------|-----------|----------|----------|---------------|----------------|
| 10.143.73.1 | 10.143.73.5 | Installed | 00:16:50 | 00:16:50 | 1700 | 1800 |
| 10.143.73.1 | 10.143.73.6 | Installed | 00:16:50 | 00:16:50 | 1700 | 1800 |

```
Total number of entries are 2
```

```
PE1#show evpn mpls xconnect tunnel label
```

```
EVPN-MPLS Network tunnel labels
```

| Remote | | Local | | Remote | | Local | | | |
|-------------|-----------|---------|---------|--------------|--------------|----------|----------|-----|--|
| Destination | Status | VPWS-ID | VPWS-ID | Network-Intf | Tunnel-Label | MC-Label | UC-Label | MC- | |
| Label | UC-Label | | | | | | | | |
| 10.143.73.5 | Installed | 1700 | 1800 | xe3 | 25601 | -- | 402 | - | |
| - 52 | | | | | | | | | |
| 10.143.73.6 | Installed | 1700 | 1800 | xe3 | 25600 | -- | 402 | - | |
| - 401 | | | | | | | | | |

```
Total number of entries are 2
```

```
PE1#show evpn mpls xconnect id 1700
```

```
EVPN-MPLS Xconnect Info
```

```
=====
```

```
AC-AC: Local-Cross-connect
```

```
AC-NW: Cross-connect to Network
```

```
AC-UP: Access-port is up
```

```
AC-DN: Access-port is down
```

```
NW-UP: Network is up
```

```
NW-DN: Network is down
```

```
NW-SET: Network and AC both are up
```

| Local | | | Remote | | Connection-Details | | |
|--------|----------|------|-----------|----------|-------------------------------|-----|--|
| VPN-ID | EVI-Name | MTU | VPN-ID | Source | Destination | PE- | |
| IP | MTU | Type | NW-Status | | | | |
| 1700 | ---- | 1500 | 1800 | xe6.1700 | 00:00:00:aa:aa:bb:bc:00:00:00 | | |

```
10.143.73.5      1500  AC-NW  NW-SET
```

```
10.143.73.6
```

```
1500  ----  ----
```

```
Total number of entries are 1
```

PE2: ELAN or EVPN-VPLS

```
PE2#show evpn mpls tunnel
```

```
EVPN-MPLS Network tunnel Entries
```

| Source | Destination | Status | Up/Down | Update | evpn-id |
|-------------|-------------|-----------|----------|----------|---------|
| 10.143.73.5 | 10.143.73.6 | Installed | 00:24:41 | 00:24:41 | 3000 |
| 10.143.73.5 | 10.143.73.1 | Installed | 00:24:38 | 00:24:38 | 3000 |

```
Total number of entries are 2
```

```
PE2#show evpn mpls tunnel label
```

```
EVPN-MPLS Network tunnel labels
```

| Destination Label | Status | evpn-id | Network-Intf | Tunnel-Label | Local | | Remote | |
|----------------------|-----------|---------|--------------|--------------|----------|----------|----------|----------|
| | | | | | MC-Label | UC-Label | MC-Label | UC-Label |
| 10.143.73.6 | Installed | 3000 | xe5 | 24965 | 753 | 402 | 654 | 53 |
| 10.143.73.1 | Installed | 3000 | xe5 | 25604 | 753 | 402 | 753 | 403 |

```
Total number of entries are 2
```

```
PE2#show evpn mpls id 3000
```

```
EVPN-MPLS Information
```

```
=====
```

```
Codes: NW - Network Port
       AC - Access Port
       (u) - Untagged
```

| VPN-ID Addr | EVI-Name Dst-Addr | EVI-Type | Type | Interface | ESI | VLAN | DF-Status | Src- |
|----------------|----------------------|----------|-------------|-----------|-------------------------------|------|-----------|------|
| 3000 | ---- | L2 | NW | ---- | ---- | ---- | ---- | ---- |
| | 10.143.73.5 | | 10.143.73.6 | | | | | |
| 3000 | ---- | L2 | NW | ---- | ---- | ---- | ---- | ---- |
| | 10.143.73.5 | | 10.143.73.1 | | | | | |
| 3000 | ---- | -- | AC | po90.300 | 00:00:00:aa:aa:bb:bc:00:00:00 | ---- | DF | ---- |

```
Total number of entries are 3
```

```
PE2#show evpn mpls mac-table
```

```
=====
```

| EVPN MPLS MAC Entries | | | | | | |
|-----------------------|-----------|--------|--------------|--------------------|-------------------------------|---------|
| VNID Ip/ESI | Interface | VlanId | Inner-VlanId | Mac-Addr Status | VTEP- AccessPortDesc | |
| 3000 | ---- | ---- | ---- | 0211.2000.03e8 | 10.143.73.1 | Dynamic |
| Remote | ----- | ----- | ----- | b86a.97cd.6a3d | 10.143.73.1 | Dynamic |
| 3000 | ---- | ---- | ---- | 0224.2000.03e8 | 00:00:00:aa:aa:bb:bc:00:00:00 | Dynamic |
| Remote | ----- | ----- | ----- | b86a.97d2.53bb | 00:00:00:aa:aa:bb:bc:00:00:00 | Dynamic |
| 3000 | po90.300 | ---- | ---- | b86a.97d2.53bb | 00:00:00:aa:aa:bb:bc:00:00:00 | Dynamic |
| Local | ----- | ----- | ----- | | | |

PE2: ELINE or EVPN-VPWS

```
PE2#show evpn mpls xconnect id 1800
```

```
EVPN-MPLS Xconnect Info
```

```
=====
```

```
AC-AC: Local-Cross-connect
```

```
AC-NW: Cross-connect to Network
```

```
AC-UP: Access-port is up
```

```
AC-DN: Access-port is down
```

```
NW-UP: Network is up
```

```
NW-DN: Network is down
```

```
NW-SET: Network and AC both are up
```

| Local | | | Remote | | Connection-Details | | |
|-------------|----------|-------|-----------|-----------|--------------------|-----------------------|--|
| VPN-ID | EVI-Name | MTU | VPN-ID | Source | Destination | PE- | |
| IP | MTU | Type | NW-Status | | | | |
| 1800 | ---- | 1500 | 1700 | po90.1700 | --- | Single Homed Port --- | |
| 10.143.73.1 | 1500 | AC-NW | NW-SET | | | | |

```
Total number of entries are 1
```

```
PE2#show evpn mpls xconnect tunnel
```

```
EVPN-MPLS Network tunnel Entries
```

| Source | Destination | Status | Up/Down | Update | local-evpn-id | remote- |
|-------------|-------------|-----------|----------|----------|---------------|---------|
| evpn-id | | | | | | |
| 10.143.73.5 | 10.143.73.1 | Installed | 00:50:18 | 00:50:18 | 1800 | 1700 |

```
Total number of entries are 1
```

```
PE2#show evpn mpls xconnect
```

```
EVPN-MPLS Xconnect Info
```

```
=====
```

```
AC-AC: Local-Cross-connect
```

```
AC-NW: Cross-connect to Network
```

```
AC-UP: Access-port is up
```

```
AC-DN: Access-port is down
```

```
NW-UP: Network is up
```

```
NW-DN: Network is down
```

```
NW-SET: Network and AC both are up
```

| Local | | | Remote | | Connection-Details | | |
|-------------|----------|-------|-----------|-----------|--------------------|-----------------------|--|
| VPN-ID | EVI-Name | MTU | VPN-ID | Source | Destination | PE- | |
| IP | MTU | Type | NW-Status | | | | |
| 1800 | ---- | 1500 | 1700 | po90.1700 | --- | Single Homed Port --- | |
| 10.143.73.1 | 1500 | AC-NW | NW-SET | | | | |

```
Total number of entries are 1
```

```
PE2#show evpn mpls xconnect tunnel label
```

```
EVPN-MPLS Network tunnel labels
```

| Remote | | | Local | Remote | | Local | | |
|-------------|-----------|---------|---------|--------------|--------------|----------|----------|-----|
| Destination | Status | VPWS-ID | VPWS-ID | Network-Intf | Tunnel-Label | MC-Label | UC-Label | MC- |
| Label | UC-Label | | | | | | | |
| 10.143.73.1 | Installed | 1800 | 1700 | xe5 | 25604 | -- | 401 | - |
| - | 402 | | | | | | | |

Total number of entries are 1

PE3: ELAN or EVPN-VPLS:

```
PE3#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
```

| Source | Destination | Status | Up/Down | Update | evpn-id |
|-------------|-------------|-----------|----------|----------|---------|
| 10.143.73.6 | 10.143.73.5 | Installed | 00:22:11 | 00:22:11 | 3000 |
| 10.143.73.6 | 10.143.73.1 | Installed | 00:22:11 | 00:22:11 | 3000 |

Total number of entries are 2

```
PE3#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels
```

| Destination Label | Status | evpn-id | Network-Intf | Tunnel-Label | Local | | Remote | |
|----------------------|-----------|---------|--------------|--------------|----------|----------|----------|-----|
| | | | | | MC-Label | UC-Label | MC-Label | UC- |
| 10.143.73.5 | Installed | 3000 | po92 | 24962 | 654 | 53 | 753 | 40 |
| 10.143.73.1 | Installed | 3000 | po92 | 24964 | 654 | 53 | 753 | 40 |

Total number of entries are 2

```
PE3#show evpn mpls id 3000
EVPN-MPLS Information
```

```
Codes: NW - Network Port
        AC - Access Port
        (u) - Untagged
```

| VPN-ID | EVI-Name | EVI-Type | Type | Interface | ESI | VLAN | DF-Status | Src- |
|--------|-------------|----------|-------------|-----------|-------------------------------|------|-----------|------|
| Addr | Dst-Addr | | | | | | | |
| 3000 | ---- | L2 | NW | ---- | ---- | ---- | ---- | ---- |
| | 10.143.73.5 | | 10.143.73.6 | | | | | |
| 3000 | ---- | L2 | NW | ---- | ---- | ---- | ---- | ---- |
| | 10.143.73.5 | | 10.143.73.1 | | | | | |
| 3000 | ---- | -- | AC | po90.300 | 00:00:00:aa:aa:bb:bc:00:00:00 | ---- | NON-DF | ---- |

Total number of entries are 3

```
PE3#show evpn mpls mac-table
```

```
=====
EVPN MPLS MAC Entries
=====
```

| VNID | Interface | VlanId | Inner-VlanId | Mac-Addr | VTEP- |
|--------|-----------|--------|--------------|----------------|-------------------------------|
| Ip/ESI | Type | | Status | AccessPortDesc | |
| 3000 | ---- | ---- | ---- | 0211.2000.03e8 | 10.143.73.1 |
| Remote | ----- | ----- | | | Dynamic |
| 3000 | ---- | ---- | ---- | b86a.97cd.6a3d | 10.143.73.1 |
| Remote | ----- | ----- | | | Dynamic |
| 3000 | ---- | ---- | ---- | 0224.2000.03e8 | 00:00:00:aa:aa:bb:bc:00:00:00 |
| Remote | ----- | ----- | | | Dynamic |
| 3000 | po90.300 | ---- | ---- | b86a.97d2.53bb | 00:00:00:aa:aa:bb:bc:00:00:00 |
| Local | ----- | ----- | | | Dynamic |

PE3: ELINE or EVPN-VPWS

```
PE3#show evpn mpls xconnect id 1800
```

```
EVPN-MPLS Xconnect Info
```

```
=====
```

```
AC-AC: Local-Cross-connect
```

```
AC-NW: Cross-connect to Network
```

```
AC-UP: Access-port is up
```

```
AC-DN: Access-port is down
```

```
NW-UP: Network is up
```

```
NW-DN: Network is down
```

```
NW-SET: Network and AC both are up
```

| Local | | | Remote | | Connection-Details | | |
|-------------|----------|-------|-----------|-----------|--------------------|-----------------------|--|
| VPN-ID | EVI-Name | MTU | VPN-ID | Source | Destination | PE- | |
| IP | MTU | Type | NW-Status | | | | |
| 1800 | ---- | 1500 | 1700 | po90.1700 | --- | Single Homed Port --- | |
| 10.143.73.1 | 1500 | AC-NW | NW-SET | | | | |

```
Total number of entries are 1
```

```
PE3#show evpn mpls xconnect tunnel
```

```
EVPN-MPLS Network tunnel Entries
```

| Source | Destination | Status | Up/Down | Update | local-evpn-id | remote-evpn-id |
|--------|-------------|--------|---------|--------|---------------|----------------|
|--------|-------------|--------|---------|--------|---------------|----------------|

| | | | | | | |
|-------------|-------------|-----------|----------|----------|------|------|
| 10.143.73.6 | 10.143.73.1 | Installed | 00:23:18 | 00:23:18 | 1800 | 1700 |
|-------------|-------------|-----------|----------|----------|------|------|

```
Total number of entries are 1
```

```
PE3#show evpn mpls xconnect tunnel label
```

```
EVPN-MPLS Network tunnel labels
```

| Remote | | Local | | Remote | | Local | |
|-------------|-----------|---------|---------|--------------|--------------|----------|----------|
| Destination | Status | VPWS-ID | VPWS-ID | Network-Intf | Tunnel-Label | MC-Label | UC-Label |
| Label | UC-Label | | | | | | |
| 10.143.73.1 | Installed | 1800 | 1700 | po92 | 24964 | -- | |
| 52 | -- | 402 | | | | | |

```
Total number of entries are 1
```

```
PE3#show evpn mpls xconnect
```

```
EVPN-MPLS Xconnect Info
```

```
=====
```

```
AC-AC: Local-Cross-connect
```

```
AC-NW: Cross-connect to Network
```

```
AC-UP: Access-port is up
```

```
AC-DN: Access-port is down
```

```
NW-UP: Network is up
```

```
NW-DN: Network is down
```

```
NW-SET: Network and AC both are up
```

| Local | | | Remote | | Connection-Details | | |
|-------------|----------|-------|-----------|-----------|--------------------|-----------------------|--|
| VPN-ID | EVI-Name | MTU | VPN-ID | Source | Destination | PE- | |
| IP | MTU | Type | NW-Status | | | | |
| 1800 | ---- | 1500 | 1700 | po90.1700 | --- | Single Homed Port --- | |
| 10.143.73.1 | 1500 | AC-NW | NW-SET | | | | |

Total number of entries are 1

EVPN-MPLS L2CP Tunneling

EVPN-MPLS services shall support transparent L2 control plane protocol tunneling via network tunnels between CE's, via Sub-ifp framework support. The protocols which are tunneled are controlled by configuration at parent Interface (of A/C ports). The below list of L2 control plane BPDUs shall be transparently tunneled across EVPN-MPLS networks based on egress tunnels when Tunnel action is configured.

- dot1x
- efm
- elmi
- lldp
- xSTP
- lacp

In case of action "DISCARD" and "PEER" appropriate behavior should see.

- PEER - the corresponding L2 control packet shall be uplifted to the CPU/control plane for processing.
- DISCARD - the corresponding L2 control packets are dropped at node.

Default behavior is PEER.

Topology

For topology, refer [Figure 48](#) and [Figure 50](#).



Note: L2CP should be enabled in all PEs access interface to have end to end traffic.

L2CP Configurations

| | |
|---|--|
| (config)#in xe1 | Entering into interface level of access side interface |
| (config-if)#l2protocol <protocol> tunnel | Enabling tunnel for the L2CP protocol |
| (config-if)#l2protocol <protocol> peer | Enabling peer for the L2CP protocol |
| (config-if)#l2protocol <protocol> discard | Enabling discard for the L2CP protocol |

Validation

```

PE2#show l2protocol processing interface <interface>
Bridge   Interface Name  Protocol    Processing Status  Hardware Status
=====  =====
-        xe8             stp         Tunnel             Peer
-        xe8             lacp        Tunnel             Peer
-        xe8             dot1x       Peer               Peer
-        xe8             lldp        Peer               Peer
-        xe8             efm         Discard            Discard
-        xe8             elmi        Discard            Discard
-        xe8             synce       Discard            Discard
PE3#show l2protocol interface pol counters
Interface pol
Peer           : lacp           : 94

```

```

Peer          : stp          : 298
Peer          : elmi         : 172
Peer          : dot1x        : 172
Discard       : stp          : 6558
Discard       : elmi         : 8326
Discard       : dot1x        : 9839

```



Note: Tunnel counter won't get incremented as per design.

```

PE4#show interface counters queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
* indicates monitor is active
+-----+-----+-----+-----+-----+-----+
+-----+
| Interface | Queue/Class-map | Q-Size | Tx pkts | Tx bytes | Dropped |
| pkts      | Dropped bytes   |         |          |          |         |
+-----+-----+-----+-----+-----+-----+
+-----+
cpu          reserved-mc (E)
2097152 6      522        0          0
cpu          bgp          (E)
1048576 2      145        0          0
cpu          rsvp-ldp     (E)
1048576 7      1060       0          0
cpu          bpdu         (E)
1048576 6481   4830520    5126       3815132
ge6          q0           (E)
1253376 350    258691      0          0
ge6          q6           (E)
1253376 114    10338       0          0
ge7          q0           (E)
1253376 6      2708        0          0
ge7          q6           (E)
1253376 2      186         0          0
ge8          q0           (E)
1253376 5177   3859861     0          0
xe12         q6           (E) 12517376
6          517          0          0
xe15         q0           (E) 12517376
2          1281         0          0
xe16         q0           (E) 12517376
2253        1694682      0          0

```



Note: CPU drop counters will increment for both peer and discard as per design.

EVPN-MPLS MAC Statistics

MAC Statistics feature provides statistics based on number of MAC's routes learned during ELAN and ETREE EVPN-MPLS service. Generic show command supported for both MH and SH services.

"show evpn mpls route-count" is used by admin to fetch all the three route type count (MAC-Only, MAC-IPV4, MAC-IPV6) for per VPN-ID of an EVPN MPLS service and/or route-type.

The details of routes can be fetched via "show evpn mpls nd-cache", "show evpn mpls arp-cache" and "show evpn mpls mac table".

MAC Statistics functionality is "Not Applicable" for EVPN E-LINE service as its point-2-point and does not have MAC advertisements functionality.

Topology

For topology, refer [Figure 48](#) and [Figure 50](#).

Configurations

Refer to above configurations to bring up the EVPN service for both SH and MH. No specific configuration needed to get the show command output. Based on the installed service and traffic below route count will be fetched.

Validation

```
PE1#show evpn mpls route-count
EVPN-MPLS Active route count information
=====
Max route count   : 32768
Active route count: 51
-----
  VNID      Total    MACONLY  MACIPv4  MACIPv6
-----
  602       17       7         5         5
  601       17       7         5         5
  801       17       7         5         5
  802        0       0         0         0
Total number of entries are 4
```

EVPN-MPLS Control-Word Support

A control-word is 4-byte optional field inserted between MPLS label stack and MPLS payload in data traffic as demarcation between MPLS labels and payload to distinguish a PW payload from an IP payload carried over the MPLS LSP, so that LSR or transit node which does deep packet inspection should not treat PW payload as IP payload and result in incorrect ECMP/load-sharing.

Some of salient points for support provided in OcNOS:

- control-word is applicable for both L2-EVPN (E-LAN, E-LINE etc) service data traffics and not applicable for L3-EVPN (IRB) traffic.
- control-word is based on static-CLI configurable along with any EVPN instance creation.

hardware-profile filter evpn-mpls-cw enable CLI must be configured for Control-word Support.

System requires reboot once after hardware profile filter control-word is enabled.



Note: The control-word for EVPN-MPLS is supported only in "local" static configurations and is neither sent, received, nor negotiated through BGP messaging.

Topology

For topology, refer [Figure 48](#) and [Figure 50](#).



Note: Control-word should be symmetrically enabled/disabled across all PE nodes to have end to end non-malformed traffic handling.

EVPN-MPLS Control-Word Configurations

Refer to [EVPN MPLS Single Homing \(page 1634\)](#) configurations to bring up the EVPN service for SH.

PE1: EVPN and MAC VRF Mapping

| | |
|--|---|
| (config)#hardware-profile filter evpn-mpls-cw enable | Enable hardware profile filter evpn-mpls-cw for Control-word support |
| (config)#evpn mpls id 2 xconnect target-mpls-id 252 control-word | Configure the EVPN-VPWS identifier with source identifier 2 and target identifier 252 along with control-word |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |
| (config)#evpn mpls id 1001 control-word | Configure the EVPN-VPLS identifier with identifier 1001 along with control-word |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-VPLS identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |

PE2: EVPN and VRF Mapping

| | |
|--|--|
| (config)#hardware-profile filter evpn-mpls-cw enable | Enable hardware profile filter evpn-mpls-cw for Control-word support |
| (config)#evpn mpls id 252 xconnect target-mpls-id 2 control-word | Configure the EVPN-ELINE identifier with source identifier 252 and target identifier 2 along with control-word |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-ELINE identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |
| (config)#evpn mpls id 1001 control-word | Configure the EVPN-ELAN identifier with identifier 1001 along with control-word |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-ELAN identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |

Refer to [EVPN MPLS Multihoming \(page 1662\)](#) configurations to bring up the EVPN service for MH.

PE1: EVPN and VRF Mapping

| | |
|---|---|
| (config)#hardware-profile filter evpn-mpls-cw | Enable hardware profile filter evpn-mpls-cw for |
|---|---|

| | |
|--|---|
| enable | Control-word support |
| (config)#evpn mpls id 1700 xconnect target-mpls-id 1800 control-word | Configure the EVPN-VPWS identifier with source identifier 1700 and target identifier 1800 along with control-word |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |
| (config)#evpn mpls id 3000 control-word | Configure the EVPN-VPLS identifier with identifier 3000 along with control-word |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-VPLS identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: EVPN and MAC VRF Mapping

| | |
|--|---|
| (config)#hardware-profile filter evpn-mpls-cw enable | Enable hardware profile filter evpn-mpls-cw for Control-word support |
| (config)#evpn mpls id 1800 xconnect target-mpls-id 1700 control-word | Configure the EVPN-VPWS identifier with source identifier 1800 and target identifier 1700 along with control-word |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "test" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| (config)#evpn mpls id 3000 control-word | Configure the EVPN-VPLS identifier with identifier 3000 along with control-word |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-VPLS identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |

PE3: EVPN and MAC VRF Mapping

| | |
|--|---|
| (config)#hardware-profile filter evpn-mpls-cw enable | Enable hardware profile filter evpn-mpls-cw for Control-word support |
| (config)#evpn mpls id 1800 xconnect target-mpls-id 1700 control-word | Configure the EVPN-VPWS identifier with source identifier 1800 and target identifier 1800 along with control-word |

| | |
|---|---|
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "test" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |
| (config)#evpn mpls id 3000 control-word | Configure the EVPN-VPLS identifier with identifier 3000 along with control-word |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Mapping vrf "vpls1001" to EVPN-VPLS identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |

Validation

show running-config evpn mpls should display control-word configuration along EVI instance. Traffic validations are performed using packet capture to observe 4-byte control word present in data traffic.

EVPN-MPLS RSVP Multipath

EVPN-MPLS services with RSVP multipath as a transport provides, Multiple RSVP tunnels which are grouped in ECMP/Multipath group to do load sharing for better optimized and efficient usages of multiple RSVP trunks towards the same Destination.

At ingress node the traffic is load-balanced based on the configured hash (L3 src/dest IP/port, L2 src/dst mac, or in-label if intermediate Autonomous segments). Each LSP path within multipath group can provide individual protection for each path (facility, 1-to-1, secondary).

Each tunnel path in multipath group cost may vary (can have different hop, with consideration for load-balancing the traffic).

Unicast traffic is load shared across all the available multipath member (but not mandatory always) where as in case of BUM only the Active member carries traffic. Addition/deletion should not have impact in traffic in case of unicast, and switchover time is not guaranteed <50ms in case of BUM traffic.

This feature supports across all EVPN service as both SH and MH modes and only BUM Traffic not supported for ELINE service alone as per RFC standard.

Topology

For topology, refer [Figure 48](#) and [Figure 50](#).

Multi Path Configurations

See the previous configurations to bring up the EVPN service for both SH and MH. Transport alone we need to change it from LDP/RSVP to RSVP Multipath as per below configurations.

| | |
|-------------------------|--|
| (config)#router rsvp | Creation of RSVP |
| (config-router)#int xel | Entering into interface level |
| (config-if)#enable-rsvp | Associate the configured RSVP to the exit interface. |

| | |
|---|---|
| (config-router)#int xe2 | Creation into interface level |
| (config-if)#enable-rsvp | Associate the configured RSVP to the exit interface |
| (Config)#rsvp-multipath MP1 | Create an RSVP multipath group 1 |
| (config-multipath)#description To-PE3-SH-ELINE-ELAN-ETREE | RSVP multipath specific description |
| (config-multipath)#to 7.7.7.7 | RSVP multipath destination prefix |
| (config)# rsvp-path P1_PE3 | Path for Resource Reservation Protocol (RSVP) with Name |
| (config-path)# 10.1.24.2 loose | Configure this explicit route path as a loose or strict hop |
| (config-path)#exit | Exit Path mode |
| (config)# rsvp-path P2_PE3 | Path for Resource Reservation Protocol (RSVP) with Name |
| (config-path)# 10.1.34.2 loose | Configure this explicit route path as a loose or strict hop |
| (config-path)#exit | Exit Path mode |
| (config)#rsvp-trunk T1_PE3 | Create an RSVP trunk T1 and enter the Trunk mode |
| (config-trunk)#primary path P1_PE3 | Specify an RSVP path to be used |
| (config-trunk)#to 7.7.7.7 | Specify the IPv4 egress (destination point) for the LSP |
| (config-trunk)#multipath-group MP1 | Associating the MP member to the RSVP trunk |
| (config)#rsvp-trunk T2_PE3 | Create an RSVP trunk T1 and enter the Trunk mode |
| (config-trunk)#primary path P2_PE3 | Specify an RSVP path to be used |
| (config-trunk)#to 7.7.7.7 | Specify the IPv4 egress (destination point) for the LSP |
| (config-trunk)#multipath-group MP1 | Associating the MP member to the RSVP trunk |
| (config-trunk)#exit | Exit trunk mode |
| (config)#commit | Commit candidate configuration to be running configuration |

Validation

Below command shows RSVP multipath group info per VPN-ID/destination.

```

PE3#show evpn mpls tunnel label
EVPN-MPLS Network tunnel labels
(*) in Policy - tunnel-policy inherited from mac-vrf
=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====
VP-Multipath      Underlay
Destination      Status   VPN-ID   Policy   Local      Remote      RS
Name             NHLFE-ix NW-Intf  NW-Label MC-Label   UC-Label   MC-Label   UC-Label   Grp-
=====+=====+=====+=====+=====+=====+=====+=====+=====
=====+=====+=====+=====
7.7.7.7           Installed 1601    --      --      19      MP2      10      NA      NA
642              20      642
7.7.7.7           Installed 601     --      --

```

| | | | | | | | | |
|---------|-----------|------|-----|-----|-----|----|----|----|
| | 640 | 17 | 640 | 17 | MP2 | 10 | NA | NA |
| 7.7.7.7 | Installed | 1801 | -- | | | | | |
| | 643 | 21 | 643 | 20 | MP2 | 10 | NA | NA |
| 7.7.7.7 | Installed | 801 | -- | | | | | |
| | 641 | 18 | 641 | 18 | MP2 | 10 | NA | NA |
| 7.7.7.7 | Installed | 1601 | -- | | | | | |
| | 642 | 20 | 18 | 643 | MP2 | 19 | NA | NA |
| 7.7.7.7 | Installed | 601 | -- | | | | | |
| | 640 | 17 | 16 | 640 | MP2 | 19 | NA | NA |
| 7.7.7.7 | Installed | 1801 | -- | | | | | |
| | 643 | 21 | 19 | 644 | MP1 | 19 | NA | NA |
| 7.7.7.7 | Installed | 1601 | -- | | | | | |
| | 642 | 20 | 640 | 17 | MP1 | 12 | NA | NA |
| 7.7.7.7 | Installed | 1801 | -- | | | | | |
| | 643 | 21 | 641 | 18 | MP1 | 12 | NA | NA |

Total number of entries are 9

```
PE3#show evpn mpls xc tunnel label
```

EVPN-MPLS Network tunnel labels

```
(*) in Policy - tunnel-policy inherited from mac-vrf
```

=====+=====+=====+=====+=====+=====+=====+=====+=====

=====+=====+=====

| Multipath | | Underlay | Local | Remote | | | Local | Remote | RSVP- |
|-------------|-----------|----------|----------|--------|--|----------|----------|--------|-------|
| Destination | Status | VPWS-ID | VPWS-ID | Policy | | UC-Label | UC-Label | Grp- | |
| Name | NHLFE-ix | NW-Intf | NW-Label | | | | | | |
| 7.7.7.7 | Installed | 501 | 1 | -- | | | | | |
| 16 | 16 | MP2 | 13 | NA | | NA | | | |
| 7.7.7.7 | Installed | 1501 | 1001 | -- | | | | | |
| 19 | 16 | MP1 | 2 | NA | | NA | | | |

Total number of entries are 2

Below show command gives info about which MP member carries traffic.

```
PE3#show mpls counters rsvp multipath-name MP2
```

Tunnel-id 5008 Extended Tunnel-ID 3.3.3.3Egress 7.7.7.7

```
lsp-name : T1 PE3-Primary [Ingress]
```

```
lsp-ingress : 3.3.3.3      lsp-id : 2208
```

```

Rx pkts : n/a
Rx bytes : n/a

```

Tx pkts : 17210 Tx bytes : 80077843

Tunnel-id 5009 Extended Tunnel-ID 7.7.7.7 Egress 3.3.3.3

```
lsp-name : T2 PE3-Primary [Ingress]
```

```
lsp-ingress : 3.3.3.3      lsp-id : 2209
```

Rx pkts : n/a Rx bytes : n/a

Tx pkts : 16904 Tx bytes : 78376789

Below command gives ingress and egress evpn mpls counters during bi-directional traffic for BUM and unicast traffic per VPN ID.

```
PE3#show evpn mpls counters network ingress
```

| | | |
|--------|-----|---------|
| VPN-ID | BUM | Unicast |
|--------|-----|---------|

| | RX (pkts) | RX (pkts) |
|-----|-----------|-----------|
| 1 | 1000 | 1000 |
| 2 | 1000 | 1000 |
| 3 | 1000 | 1000 |
| 4 | 1000 | 1000 |
| 5 | 1000 | 1000 |
| 6 | 1000 | 1000 |
| 7 | 1000 | 1000 |
| 8 | 1000 | 1000 |
| 9 | 1000 | 1000 |
| 10 | 1000 | 1000 |
| 11 | 1000 | 1000 |
| 12 | 1000 | 1000 |
| 13 | 1000 | 1000 |
| 14 | 1000 | 1000 |
| 15 | 1000 | 1000 |
| 16 | 1000 | 1000 |
| 17 | 1000 | 1000 |
| 18 | 1000 | 1000 |
| 19 | 1000 | 1000 |
| 20 | 1000 | 1000 |
| 21 | 1000 | 1000 |
| 22 | 1000 | 1000 |
| 23 | 1000 | 1000 |
| 24 | 1000 | 1000 |
| 25 | 1000 | 1000 |
| 26 | 1000 | 1000 |
| 27 | 1000 | 1000 |
| 28 | 1000 | 1000 |
| 29 | 1000 | 1000 |
| 30 | 1000 | 1000 |
| 31 | 1000 | 1000 |
| 32 | 1000 | 1000 |
| 33 | 1000 | 1000 |
| 34 | 1000 | 1000 |
| 35 | 1000 | 1000 |
| 36 | 1000 | 1000 |
| 37 | 1000 | 1000 |
| 38 | 1000 | 1000 |
| 39 | 1000 | 1000 |
| 40 | 1000 | 1000 |
| 41 | 1000 | 1000 |
| 42 | 1000 | 1000 |
| 43 | 1000 | 1000 |
| 44 | 1000 | 1000 |
| 45 | 1000 | 1000 |
| 46 | 1000 | 1000 |
| 47 | 1000 | 1000 |
| 48 | 1000 | 1000 |
| 49 | 1000 | 1000 |
| 50 | 1000 | 1000 |
| 51 | 1000 | 1000 |
| 52 | 1000 | 1000 |
| 53 | 1000 | 1000 |
| 54 | 1000 | 1000 |
| 55 | 1000 | 1000 |
| 56 | 1000 | 1000 |
| 57 | 1000 | 1000 |
| 58 | 1000 | 1000 |
| 59 | 1000 | 1000 |
| 60 | 1000 | 1000 |
| 61 | 1000 | 1000 |
| 62 | 1000 | 1000 |
| 63 | 1000 | 1000 |
| 64 | 1000 | 1000 |
| 65 | 1000 | 1000 |
| 66 | 1000 | 1000 |
| 67 | 1000 | 1000 |
| 68 | 1000 | 1000 |
| 69 | 1000 | 1000 |
| 70 | 1000 | 1000 |
| 71 | 1000 | 1000 |
| 72 | 1000 | 1000 |
| 73 | 1000 | 1000 |
| 74 | 1000 | 1000 |
| 75 | 1000 | 1000 |
| 76 | 1000 | 1000 |
| 77 | 1000 | 1000 |
| 78 | 1000 | 1000 |
| 79 | 1000 | 1000 |
| 80 | 1000 | 1000 |
| 81 | 1000 | 1000 |
| 82 | 1000 | 1000 |
| 83 | 1000 | 1000 |
| 84 | 1000 | 1000 |
| 85 | 1000 | 1000 |
| 86 | 1000 | 1000 |
| 87 | 1000 | 1000 |
| 88 | 1000 | 1000 |
| 89 | 1000 | 1000 |
| 90 | 1000 | 1000 |
| 91 | 1000 | 1000 |
| 92 | 1000 | 1000 |
| 93 | 1000 | 1000 |
| 94 | 1000 | 1000 |
| 95 | 1000 | 1000 |
| 96 | 1000 | 1000 |
| 97 | 1000 | 1000 |
| 98 | 1000 | 1000 |
| 99 | 1000 | 1000 |
| 100 | 1000 | 1000 |

-----+

501 0 0

| | | |
|-----|---|---|
| 501 | 0 | 0 |
| 601 | 0 | 0 |

| | | |
|-----|---|---|
| 001 | 0 | 0 |
| 801 | 0 | 0 |

| | | |
|------|---|---|
| 801 | 0 | 0 |
| 1501 | 0 | 0 |

| | | |
|------|--------|---|
| 1501 | 0 | 0 |
| 1601 | 131654 | 0 |

| | | |
|------|--------|---|
| 1801 | 131654 | 0 |
| 1801 | 131682 | 0 |

| | | |
|------|--------|---|
| 1801 | 131682 | 0 |
| REF# | | |

PE3#
PE3#

```
PE3#show evpn mpls counters network egress
```

| VPN-ID | DESTINATION PEER | BUM TX (pkts) | Unicast TX (pkts) |
|--------|---------------------|------------------|----------------------|
| 1601 | 3.3.3.3 | 264177 | 0 |
| 501 | 3.3.3.3 | 0 | 0 |
| 601 | 3.3.3.3 | 0 | 0 |
| 1801 | 3.3.3.3 | 264312 | 0 |
| 801 | 3.3.3.3 | 0 | 0 |
| 1601 | 8.8.8.8 | 264439 | 0 |
| 601 | 8.8.8.8 | 0 | 0 |
| 1801 | 8.8.8.8 | 264440 | 0 |
| 1601 | 4.4.4.4 | 264503 | 0 |
| 1501 | 4.4.4.4 | 0 | 0 |
| 1801 | 4.4.4.4 | 264530 | 0 |

EVPN-ELINE CFM Single Homing

CFM (according to IEEE 802.1ag 2007) offers capabilities that are useful for detecting, verifying, and isolating connectivity failures in Virtual Bridged Local Area Networks. This is achieved through Continuity Check, Loop Back, and Link Trace protocols. These capabilities can be utilized in networks operated by multiple independent organizations, each with limited access to each other's equipment.

Typically, the network administrator is informed about the failure in the connection based on the receipt of Continuity Check Messages or by the user. The administrator can then initiate Loop Back or Link Trace as needed to quickly identify and isolate the fault condition.


The CFM information is conveyed in Protocol frames, which are called CFM Protocol Data Units (CFM PDUs). The CFM PDUs contain the necessary control and status information used to detect, verify and isolate faults. They also contain information for path discovery in CFM-enabled links.

Currently, this is only supported for EVPN-ELINE Single home and up MEP service on both Q1 and Q2 platforms.

Topology

The diagram depicts the Single Homed topology for the EVPN MPLS configuration examples that follow.

Figure 52. EVPN-ELINE CFM Single Homing configuration



Prerequisite


Configure the below hardware-profile commands related to CFM in configuration mode and reboot the nodes.

```
hardware-profile filter cfm-domain-name-str enable
hardware-profile statistics cfm-ccm enable
```


PE1: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.1/32 secondary | Configure IP address on loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE1: Global EVPN MPLS Command

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#commit | Commit candidate configuration to be running configuration |
| |  Note: Reload is required after Enabling or Disabling EVPN MPLS. |
| (config)#evpn mpls vtep-ip-global 10.143.73.1 | Configuring VTEP global IP to loopback IP |

PE1: Global LDP

| | |
|---|---|
| (config)#router ldp | Enter the Router LDP mode |
| PE1(config-router)#router-id 10.143.73.1 | Set the router ID to IP address 10.143.73.1 |
| PE1(config-router)#transport-address ipv4 10.143.73.1 0 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run |
| |  Note: It is preferable to use the loopback address as the transport address. |
| PE1(config-router)#targeted-peer ipv4 10.143.73.3 | Configure targeted peer |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE1: Interface Configuration Network Side

| | |
|--|---|
| (config)#interface xe3 | Enter the Interface mode for xe3 |
| (config-if)#ip address 10.255.128.2/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |

| | |
|------------------|--|
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |



Note: For RSVP Configuration refer [RSVP-TE Configuration \(page 643\)](#).

PE1: OSPF Configuration

| | |
|--|--|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 10.143.73.1 | Router-ID configurations |
| (config-router)#network 10.143.73.1/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.2/31 area 0.0.0.0 | Advertise xe3 network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE1: BGP Configuration

| | |
|---|--|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.3 remote-as 65010 | Configuring PE2 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 10.143.73.3 update-source lo | Source of routing updates as loopback |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.3 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config-router)#commit | Commit candidate configuration to be running configuration |

PE1: MAC VRF Configuration

| | |
|------------------------------------|--|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.1:2 | Configuring Route-Distinguisher value 10.143.73.1:2 |
| (config-vrf)#route-target both 2:2 | Configuring import and export value as 2:2 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE1: EVPN and VRF Mapping

| | |
|---|---|
| (config)#evpn mpls id 2 xconnect target-mpls-id 252 | Configure the EVPN-VPWS identifier with source identifier 2 and target identifier 252 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |

PE1: Access Port Configuration

| | |
|---|--|
| (config)#interface xe6 | Enter the Interface mode for xe6 |
| (config-if)#description access-side-int | Giving Interface Description |
| (config-if)#interface xe6.2 switchport | Creating L2 sub interface of physical interface xe6 |
| (config-if)#encapsulation dot1q 2 | Setting Encapsulation to dot1q with VLAN ID 2 Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if)#map vpn-id 2 | Map vpn-id 252 to interface xe2.2 (VPWS) |
| (config-access-if)#exit | Exiting out of access interface mode |
| (config-if)#commit | Commit candidate configuration to be running configuration |

PE1: CFM Configuration


| | |
|--|---|
| (config)#hardware-profile filter cfm-domain-name-str enable | Configure cfm-domain-name-str profile to enable cfm |
| (config)#ethernet cfm domain-type character-string domain-name MD-01 level 2 mip-creation none | Create CFM domain for Evpn ELine with type as character string and set mip creation to none |
| (config-ether-cfm-mpls-md)#service ma-type string ma-name S1 | Create ma type with string and set mip creation to none |
| (config-ether-cfm-mpls-ma)# evpn 2 | Configure evpn <Evpn-id> |
| (config-ether-cfm-mpls-ma)#ethernet cfm mep up mpid 8191 active true evpn 2 | Create up-mep for local evpn id 2 |
| (config-ether-cfm-mpls-ma-mep)#cc multicast state enable | Enable cc multicast |
| (config-ether-cfm-mpls-ma-mep)#exit-ether-ma-mep-mode | Exit Ethernet ma-mep-mode |
| (config-ether-cfm-mpls-ma)#mep crosscheck mpid 8000 | Configure cross check to remote mep for vlan 2 |
| (config-ether-cfm-mpls-ma)#cc interval 2 | Enable cc interval with value 2 i.e 10 milliseconds |
| (config-ether-cfm-mpls-ma)#exit-ether-ma-mode | Exit Ethernet ma mode |
| (config-ether-cfm-mpls-md)#exit | Exit Ethernet cfm mode |

| | |
|-----------------|--|
| (config)#exit | Exit Config mode |
| (config)#commit | Commit candidate configuration to be running configuration |

P: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.2/32 secondary | Configure IP address on loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |

P: Global LDP

| | |
|---|--|
| (config)#router ldp | Enter the Router LDP mode |
| PE1(config-router)#router-id 10.143.73.2 | Set the router ID to IP address 10.143.73.2 |
| PE1(config-router)#transport-address ipv4 10.143.73.2 0 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run |
| | <div>  Note: It is preferable to use the loopback address as the transport address. </div> |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

P: Interface Configuration

| | |
|---|--|
| (config)#interface xe3 | Enter the Interface mode for xe3 |
| (config-if)#ip address 10.255.128.3/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)#ip address 10.255.128.22/31 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |


P: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode |
| (config-router)#ospf router-id 10.143.73.2 | Setting the Router ID as Loopback IP |
| (config-router)#network 10.143.73.2/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.22/31 area 0.0.0.0 | Advertise xe5 network address in OSPF |
| (config-router)#network 10.255.128.3/31 area 0.0.0.0 | Advertise xe3 network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: Loopback Interface


| | |
|---|--|
| #configure terminal | Enter configuration mode |
| (config)#interface lo | Enter the Interface mode for the loopback interface |
| (config-if)#ip address 10.143.73.3/32 secondary | Configure IP address on loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: Global LDP

| | |
|---|---|
| (config)#router ldp | Enter the Router LDP mode |
| PE1(config-router)#router-id 10.143.73.3 | Set the router ID to IP address 10.143.73.3 |
| PE1(config-router)#transport-address ipv4 10.143.73.3 0 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run |
| |  Note: It is preferable to use the loopback address as the transport address. |
| PE1(config-router)#targeted-peer ipv4 10.143.73.1 | Configure targeted peer |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: Global EVPN MPLS Command

| | |
|---------------------------|--|
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#commit | Commit candidate configuration to be running configuration |

| | |
|---|--|
| |  Note: Reload is required after Enabling or Disabling EVPN MPLS. |
| (config)#evpn mpls vtep-ip-global 10.143.73.3 | Configuring VTEP global IP to loopback IP |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: Interface Configuration Network Side

| | |
|---|--|
| (config)#interface xe5 | Enter the Interface mode for xe5 |
| (config-if)#ip address 10.255.128.23/31 | Configure IP address on the interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: OSPF Configuration

| | |
|---|--|
| (config)#router ospf 100 | Enter the Router OSPF mode |
| (config-router)#ospf router-id 10.143.73.3 | Router-ID configurations |
| (config-router)#network 10.143.73.3/32 area 0.0.0.0 | Advertise loopback address in OSPF |
| (config-router)#network 10.255.128.23/31 area 0.0.0.0 | Advertise xe5 network address in OSPF |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: BGP Configuration

| | |
|---|--|
| (config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| (config-router)#neighbor 10.143.73.1 remote-as 65010 | Configuring PE1 as iBGP neighbor using its loopback IP |
| (config-router)#neighbor 10.143.73.1 update-source lo | Source of routing updates as loopback |
| (config-router)#address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)#neighbor 10.143.73.1 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exiting of Address family mode |
| (config-router)#commit | Commit candidate configuration to be running configuration |

PE2: MAC VRF Configuration

| | |
|------------------------------------|--|
| (config)#mac vrf vrf2 | Enter VRF mode |
| (config-vrf)#rd 10.143.73.3:2 | Configuring Route-Distinguisher value 10.143.73.3:2 |
| (config-vrf)#route-target both 2:2 | Configuring import and export value as 2:2 |
| (config-vrf)#exit | Exiting VRF Mode |
| (config)#commit | Commit candidate configuration to be running configuration |

PE2: EVPN and VRF Mapping

| | |
|---|--|
| (config)#evpn mpls id 252 xconnect target-mpls-id 2 | Configure the EVPN-ELINE identifier with source identifier 252 and target identifier 2 |
| (config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Mapping vrf "vrf2" to EVPN-ELINE identifier |
| (config-evpn-mpls)#commit | Commit candidate configuration to be running configuration |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode |

PE2: Access Port Configuration:

| | |
|---|--|
| (config)#interface xe2 | Enter the Interface mode for xe2 |
| (config-if)#description access-side-int | Giving Interface Description |
| (config-if)#interface xe2.2 switchport | Creating L2 sub interface of physical interface xe2 |
| (config-if)#encapsulation dot1q 2 | Setting Encapsulation to dot1q with VLAN ID 2 Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-access-if)#map vpn-id 252 | Map vpn-id 252 to interface xe2.2 (VPWS) |
| (config-access-if)#exit | Exiting out of access interface mode |
| (config-if)#commit | Commit candidate configuration to be running configuration |

PE2: CFM Configuration

| | |
|--|---|
| (config)#hardware-profile filter cfm-domain-name-str enable | Configure cfm-domain-name-str profile to enable cfm |
| (config)#ethernet cfm domain-type character-string domain-name MD-01 level 2 mip-creation none | Create CFM domain for Evpn ELine with type as character string and set mip creation to none |
| (config-ether-cfm-mpls-md)#service ma-type string ma-name S1 | Create ma type with string and set mip creation to none |

| | |
|---|--|
| (config-ether-cfm-mpls-ma)# evpn 252 | Configure evpn <Evpn-id> |
| (config-ether-cfm-mpls-ma)#ethernet cfm mep up mpid 8000 active true evpn 252 | Create up-mep for local evpn id 252 |
| (config-ether-cfm-mpls-ma-mep)#cc multicast state enable | Enable cc multicast |
| (config-ether-cfm-mpls-ma-mep)#exit-ether-ma-mep-mode | Exit Ethernet ma-mep-mode |
| (config-ether-cfm-mpls-ma)#mep crosscheck mpid 8191 | Configure cross check to remote mep for vlan 2 |
| (config-ether-cfm-mpls-ma)#cc interval 2 | Enable cc interval with value 2 i.e 10 milliseconds |
| (config-ether-cfm-mpls-ma)#exit-ether-ma-mode | Exit ethernet ma mode |
| (config-ether-cfm-mpls)#exit | Exit Ethernet cfm mode |
| (config)#exit | Exit from config mode |
| (config)#commit | Commit candidate configuration to be running configuration |

Validation

PE1

```

PE1#show evpn mpls xconnect id 2
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                               Remote      Connection-Details
=====
VPN-ID      EVI-Name      MTU  VPN-ID      Source      Destination      PE-
IP           MTU   Type   NW-Status
=====
2           ----      1500  252         xe6.2       --- Single Homed Port ---
      10.143.73.3      1500  AC-NW  NW-SET

Total number of entries are 1

PE1#show ethernet cfm errors domain MD-01

Domain Name      Level      MEPID      Defects
-----
MD-01            2          8191      .....

1. defRDICCM      2. defMACstatus  3. defRemoteCCM
4. defErrorCCM    5. defXconCCM

PE1#show ethernet cfm ma status domain MD-01 ma-name S1
MA NAME          STATUS
-----
S1               Active

PE1#show ethernet cfm maintenance-points local mep domain MD-01 ma-name S1
MPID Dir Lvl CC-Stat HW-Status CC-Intvl MAC-Address Def Port MD Name

```

```

-----
8191 Up 2 Enable Installed 10 ms e8c5.7a78.7124 F xe6.2 MD-01

PE1#show ethernet cfm maintenance-points remote domain MD-01 ma-name S1
MEPID RMEPID LEVEL Rx CCM RDI PEER-MAC TYPE
-----
8191 8000 2 Yes False b86a.97cb.6c6e Configured

PE1#show ethernet cfm maintenance-points remote mpid 8191 domain MD-01 ma-name S1
MEPID RMEPID LEVEL Rx CCM RDI PEER-MAC TYPE
-----
8191 8000 2 Yes False b86a.97cb.6c6e Configured

PE1#ping ethernet mac b86a.97cb.6c6e unicast source 8191 domain MD-01 ma-name S1
success rate is 100 (5/5)

PE1#traceroute ethernet b86a.97cb.6c6e mepid 8191 domain MD-01 ma-name S1
MP Mac Hops Relay-action Ingress/Egress Ingress/Egress action
b86a.97cb.6c6e 1 RlyHit Ingress IngOK

PE1#show ethernet cfm statistics mep 8191 domain MD-01

CFM Statistics for MEP 8191 of MD MD-01
=====
Continuity Check Messages
  CCM Sent : CCM Stats Profile Disabled
  CCM Received : CCM Stats Profile Disabled

Loop Back Messages
  LBM Sent : 5
  LBR Received(Valid) : 5
  LBR Received(Bad msdu) : 0
  LBR Received(Out-of-Seq): 0

Link Trace Messages
  LTM Sent : 1
  LTR Sent : 0
  LTR Received(Valid) : 1
  LTR Received(unexpected): 0

```

PE2

```

PE2#show evpn mpls xconnect id 252
EVPN-MPLS Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local Remote Connection-Details
=====
=====
VPN-ID EVI-Name MTU VPN-ID Source Destination PE-
IP MTU Type NW-Status
=====
=====
252 ---- 1500 2 xe2.2 --- Single Homed Port ---
10.143.73.1 1500 AC-NW NW-SET

Total number of entries are 1

```

```
PE2#show ethernet cfm errors domain MD-01
```

| Domain Name | Level | MEPID | Defects |
|-------------|-------|-------|---------|
| MD-01 | 2 | 8000 | |

```
1. defRDICCM    2. defMACstatus  3. defRemoteCCM
4. defErrorCCM  5. defXconCCM
```

```
PE3#show ethernet cfm ma status domain MD-01 ma-name S1
```

```
ma-name S1
```

| MA NAME | STATUS |
|---------|--------|
| S1 | Active |

```
PE2#show ethernet cfm maintenance-points local mep domain MD-01 ma-name S1
```

| MPID | Dir | Lvl | CC-Stat | HW-Status | CC-Intvl | MAC-Address | Def Port | MD Name |
|------|-----|-----|---------|-----------|----------|----------------|----------|-------------|
| 8000 | Up | 2 | Enable | Installed | 10 ms | b86a.97cb.6c6e | F | xe2.2 MD-01 |

```
PE3#show ethernet cfm maintenance-points remote domain MD-01 ma-name S1
```

| MEPID | RMEPID | LEVEL | Rx CCM | RDI | PEER-MAC | TYPE |
|-------|--------|-------|--------|-------|----------------|------------|
| 8000 | 8191 | 2 | Yes | False | e8c5.7a78.7124 | Configured |

```
PE2#show ethernet cfm maintenance-points remote mpid 8000 domain MD-01 ma-name S1
```

| MEPID | RMEPID | LEVEL | Rx CCM | RDI | PEER-MAC | TYPE |
|-------|--------|-------|--------|-------|----------------|------------|
| 8000 | 8191 | 2 | Yes | False | e8c5.7a78.7124 | Configured |

```
PE2#traceroute ethernet e8c5.7a78.7124 mepid 8000 domain MD-01 ma-name S1
```

| MP Mac | Hops | Relay-action | Ingress/Egress | Ingress/Egress action |
|----------------|------|--------------|----------------|-----------------------|
| e8c5.7a78.7124 | 1 | RlyHit | Ingress | IngOK |

```
PE2#show ethernet cfm statistics mep 8000 domain MD-01 ma-name S1
```

```
CFM Statistics for MEP 8000 of MD MD-01
```

```
=====
```

```
Continuity Check Messages
```

```
CCM Sent           : CCM Stats Profile Disabled
CCM Received        : CCM Stats Profile Disabled
```

```
Loop Back Messages
```

```
LBM Sent           : 5
LBR Received(Valid) : 5
LBR Received(Bad msdu) : 0
LBR Received(Out-of-Seq) : 0
```

```
Link Trace Messages
```

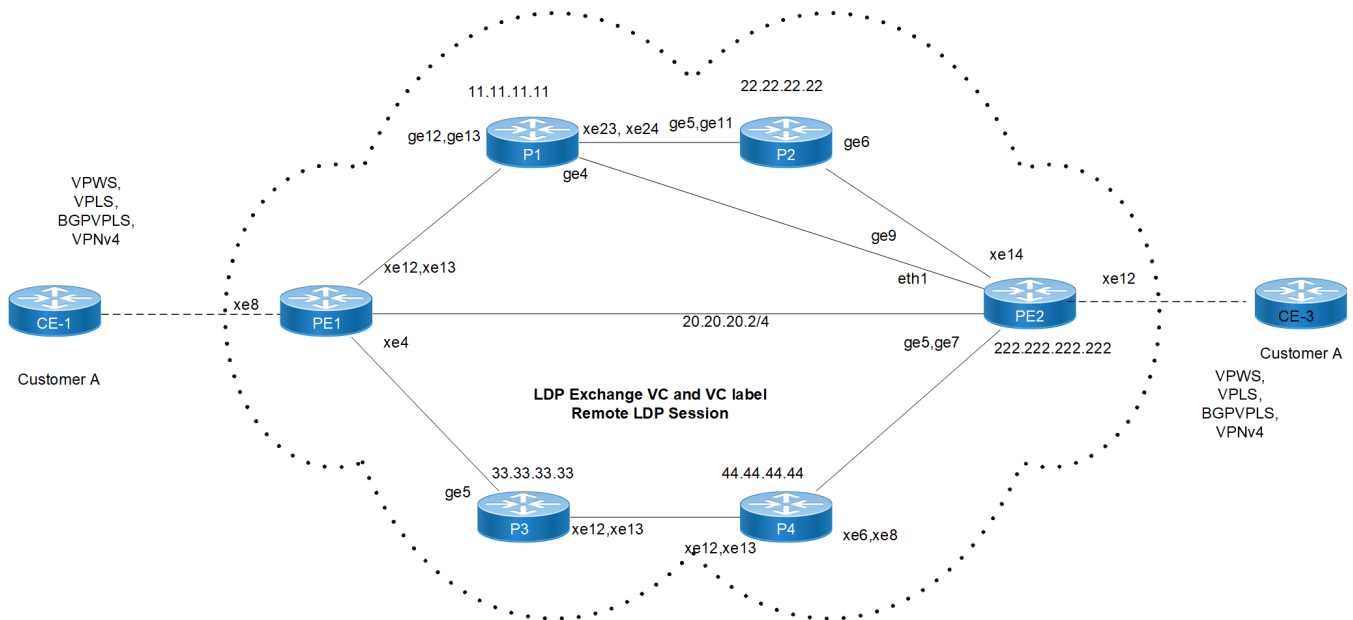
```
LTM Sent           : 1
LTR Sent           : 1
LTR Received(Valid) : 1
LTR Received(unexpected) : 0
```

LDP ECMP Configuration

This chapter contains configurations of LDP ECMP detailed tests that are used to verify the functionality of LDP ECMP (Equal-cost multipath).

Topology

Figure 53. LDP-ECMP configuration topology



Configurations

PE-1

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 111.111.111.111/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 111.111.111.111 | Configure Router ID |
| (config-router)#targeted-peer ipv4 222.222.222.222 | Configuring targeted LDP sessions to PE-2 |
| (config-router)#explicit-null | Configure explicit-null. |
| (config-router)#entropy-label-capability | Enable entropy capability in ldp |
| (config-router-targeted-peer)#exit-targeted-peer- | Exit from targeted-peer mode |

| | |
|--|---|
| mode | |
| (config-router)#transport-address ipv4 111.111.111.111 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface xe13 | Enter interface mode |
| (config-if)#speed 1g | Configure interface speed to 1g |
| (config-if)#ip address 10.0.1.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po1000 | Enter the Interface mode for po1000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface po1000.1000 | Enter interface mode for po1000.1000 |
| (config-if)#encapsulation dot1q 1000 | Configure encapsulation under a subinterface |
| (config-if)#ip address 10.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe4 | Enter interface mode |
| (config-if)#speed 1g | Configure interface speed to 1g |
| (config-if)#ip address 16.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe4.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 16.0.1.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |

| | |
|---|---|
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface xe12 | Enter interface mode |
| (config-if)#speed 1g | Configure interface speed to 1g |
| (config-if)#channel-group 1000 mode active | Moving interface to Dynamic LAG |
| (config-if)#exit | Exit interface mode |
| (config)#ip vrf l3vpnvrf300 | IP VRF config with name l3vpnvrf300 |
| (config-vrf)#rd 300:1 | Route-distinguisher value |
| (config-vrf)#route-target both 300:1 | Route target value |
| (config-vrf)#exit | Exit to config mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 111.111.111.111 | Configure ospf Router-id |
| (config-router)#network 10.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 16.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 16.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 111.111.111.111/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#router bgp 100 | Configure router bgp in AS 100 |
| (config-router)#bgp router-id 111.111.111.111 | Configure BGP router ID |
| (config-router)#neighbor 222.222.222.222 remote-as 100 | Configure neighbor in remote-as 100 |
| (config-router)#neighbor 222.222.222.222 update-source lo | Configure neighbor with update-source loopback |
| (config-router)#address-family vpnv4 unicast | Enter VPNv4 Address family mode |
| (config-router-af)#neighbor 222.222.222.222 activate | mode Activate VPNv4 neighbor |
| (config-router-af)#exit-address-family | Exit from address Fmily configuration |
| (config-router)#address-family l2vpn vpls | Enter VPLS address family mode |
| (config-router-af)#neighbor 222.222.222.222 activate | mode Activate vpls neighbor |

| | |
|--|---|
| (config-router-af) #exit-address-family | Exit from Address Family configuration |
| (config-router) #address-family vrf l3vpnvr300 | Configure VRF address family |
| (config-router-af) #redistribute connected | Redistribute connected addresses |
| (config-router-af) #exit-address-family | Exit from Address Family configuration |
| (config-router) #exit | Exit from router mode |
| (config) #mpls vplsldp100 100 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls) #signaling ldp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig) # vpls-peer 222.222.222.222 | Configuring VPLS mesh peers |
| (config-vpls-sig) #exit-signaling | Exit from VPLS signaling mode |
| (config-vpls) #exit-vpls | Exit from VPLS mode |
| (config) #mpls l2-circuit VPWS400 400 222.222.222.222 | Configuring VPWS instance with name and VPWS ID |
| (config) #mpls vpls vplsbgp200 200 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls) #signaling bgp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig) #ve-id 200 | Configure VE ID, which is mandatory for BGP VPLS, otherwise, signaling does not take place. VE ID should be unique per VPLS instance. |
| (config-vpls-sig) #exit-signaling | Exit from VPLS signaling mode |
| (config-vpls) #exit-vpls | Exit from VPLS mode |
| (config-if) #interface xe8.100 switchport | Enter sub interface mode |
| (config-if) #encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if) #access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls) #mpls-vpls vplsldp100 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls) #exit | Exit from access mode |
| (config-if) #interface xe8.400 switchport | Enter sub interface mode |
| (config-if) #encapsulation dot1q 400 | Configure encapsulation under a subinterface |
| (config-if) #access-if-vpws | Access VPWS under sub interface |
| (config-acc-if-vpws) #mpls-vpws VPWS400 | Associating the VPWS Instance to the attachment circuit interface. |
| (config-acc-if-vpws) #exit | Exit from access mode |
| (config-if) #interface xe8.200 switchport | Enter sub interface mode |
| (config-if) #split-horizon group access1 | Configure split-horizon group on sub-interface |
| (config-if) #split-horizon group access1 | Configure split-horizon group on sub-interface |
| (config-if) #encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if) #access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls) #mpls-vpls vplsbgp200 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls) #exit | Exit from access mode |

| | |
|---|--|
| (config-if)#interface xe8.300 | Enter sub interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#ip vrf forwarding l3vpnvrf300 | Attaching xe8.300 to as part of l3vpnvrf300 |
| (config-if)#ip address 110.110.110.1/24 | Configure the IP address of the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the configuration |
| (config)#end | Return to privilege mode |

P1

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#mpls ilm-ecmp | Enable ILM ECMP |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 11.11.11.11/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 11.11.11.11 | Configure Router-id |
| (config-router)#transport-address ipv4 11.11.11.11 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface ge13 | Enter interface mode |
| (config-if)#ip address 10.0.1.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po1000 | Enter the Interface mode for po1000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface po1000.1000 | Enter interface mode for po1000.1000 |
| (config-if)#encapsulation dot1q 1000 | Configure encapsulation under a subinterface |
| (config-if)#ip address 10.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |

| | |
|--|---|
| (config-if)#exit | Exit interface mode |
| (config)#interface po2000 | Enter the Interface mode for po2000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface po2000.100 | Enter interface mode for po2000.100 |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#ip address 11.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge4 | Enter interface mode |
| (config-if)#ip address 9.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 20 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge12 | Enter interface mode |
| (config-if)#channel-group 1000 mode active | Moving interface to Dynamic LAG |
| (config)#interface xe23 | Enter interface mode |
| (config-if)#channel-group 2000 mode active | Moving interface to Dynamic LAG |
| (config-if)#speed 1g | Configure speed 1g |
| (config)#interface xe24 | Enter interface mode |
| (config-if)#channel-group 2000 mode active | Moving interface to Dynamic LAG |
| (config-if)#speed 1g | Configure speed 1g |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 11.11.11.11 | Configure ospf Router-id |
| (config-router)#network 10.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 11.0.0.0/24 area 0 | Define the interface on which OSPF runs and |

| | |
|---|---|
| | associate the area ID (0) with the interface. |
| (config-router)#network 9.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 11.11.11.11/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#exit | Exit from router mode |
| (config)#commit | Commit the transaction |

P2

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 22.22.22.22/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 22.22.22.22 | Configure Router ID |
| (config-router)#transport-address ipv4 22.22.22.22 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface ge6.300 | Enter interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#ip address 15.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge6.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 15.0.1.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po2000 | Enter the Interface mode for po2000 |

| | |
|---|---|
| (config-if)#exit | Exit interface mode |
| (config)#interface po2000.100 | Enter interface mode for po2000.100 |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#ip address 11.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge5 | Enter interface mode |
| (config-if)#channel-group 2000 mode active | Moving interface to Dynamic LAG |
| (config)#interface ge11 | Enter interface mode |
| (config-if)#channel-group 2000 mode active | Moving interface to Dynamic LAG |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 22.22.22.22 | Configure ospf Router ID |
| (config-router)#network 15.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 15.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 11.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 22.22.22.22/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#exit | Exit from router mode |
| (config)#commit | Commit the transaction |

P3

| | |
|---|---|
| #configure terminal | Enter Configure mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 33.33.33.33/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |

| | |
|--|---|
| (config-router)#router-id 33.33.33.33 | Configure Router-id |
| (config-router)#transport-address ipv4 33.33.33.33 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface ge5 | Enter interface mode |
| (config-if)#ip address 16.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge5.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 16.0.1.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa4000 | Enter the Interface mode for sa4000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa4000.200 | Enter interface mode for psa4000.200 |
| (config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if)#ip address 17.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe12 | Enter interface mode |
| (config-if)# static-channel-group 4000 | Moving interface to Static LAG |
| (config)#interface xe13 | Enter interface mode |
| (config-if)# static-channel-group 4000 | Moving interface to Static LAG |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the |

| | |
|---|---|
| | Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 33.33.33.33 | Configure ospf Router-id |
| (config-router)#network 16.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 16.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 17.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 33.33.33.33/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#exit | Exit from router mode |
| (config)#commit | Commit the transaction. |

P4

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 44.44.44.44/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 44.44.44.44 | Configure Router ID |
| (config-router)#transport-address ipv4 44.44.44.44 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface sa3000 | Enter the Interface mode for sa3000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa3000 | Enter interface mode |
| (config-if)#ip address 18.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa3000.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |

| | |
|--|---|
| (config-if)#ip address 18.0.1.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa4000 | Enter the Interface mode for sa4000 |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa4000.200 | Enter interface mode for psa4000.200 |
| (config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if)#ip address 17.0.0.10/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe6 | Enter interface mode |
| config-if)#speed 1g | Configure speed 1g |
| (config-if)#static-channel-group 3000 | Moving interface to Static LAG |
| (config)#interface xe8 | Enter interface mode |
| config-if)#speed 1g | Configure speed 1g |
| (config-if)#static-channel-group 3000 | Moving interface to Static LAG |
| (config)#interface xe12 | Enter interface mode |
| (config-if)# static-channel-group 4000 | Moving interface to Static LAG |
| (config)#interface xe13 | Enter interface mode |
| (config-if)# static-channel-group 4000 | Moving interface to Static LAG |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 44.44.44.44 | Configure OSPF Router ID |
| (config-router)#network 18.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 18.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 17.0.0.0/24 area 0 | Define the interface on which OSPF runs and |

| | |
|---|---|
| | associate the area ID (0) with the interface. |
| (config-router)#network 44.44.44.44/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#exit | Exit from router mode |
| (config)#commit | Commit the transaction. |

PE2

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 222.222.222.222/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 222.222.222.222 | Configure Router-id |
| (config-router)#targeted-peer ipv4 111.111.111.111 | Configuring targeted LDP sessions to PE-2 |
| (config-router)#explicit-null | Configure explicit-null. |
| (config-router)#entropy-label-capability | Enable entropy capability in ldp |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)#transport-address ipv4 222.222.222.222 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface xe14 | Enter interface mode |
| (config-if)#speed 1g | Configure interface speed to 1g |
| (config)#interface xe14.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 15.0.1.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe14.1002 | Enter interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#ip address 15.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign ospf cost to the interface |

| | |
|--|---|
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa3000 | Enter interface mode |
| (config-if)#ip address 18.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface sa3000.1001 | Enter interface mode |
| (config-if)#encapsulation dot1q 1001 | Configure encapsulation under a subinterface |
| (config-if)#ip address 18.0.1.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 10 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface ge5 | Enter interface mode |
| (config-if)# static-channel-group 3000 | Moving interface to static LAG |
| (config-if)#interface ge7 | Enter interface mode |
| (config-if)# static-channel-group 3000 | Moving interface to static LAG |
| (config-if)#exit | Exit interface mode |
| (config)#interface ge9 | Enter interface mode |
| (config-if)#ip address 9.0.0.20/24 | Configure IP address on interface |
| (config-if)#ip ospf cost 20 | Assign OSPF cost to the interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#ip vrf l3vpnvrf300 | IP VRF config with name l3vpnvrf300 |
| (config-vrf)#rd 300:1 | Route-distinguisher value |
| (config-vrf)#route-target both 300:1 | Route target value |
| (config-vrf)#exit | Exit to config mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a |

| | |
|---|---|
| | unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 222.222.222.222 | Configure ospf Router-id |
| (config-router)#network 9.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 15.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 15.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 18.0.0.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 18.0.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 222.222.222.222/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#router bgp 100 | Configure router bgp in AS 100 |
| (config-router)#bgp router-id 222.222.222.222 | Configure BGP router ID |
| (config-router)#neighbor 111.111.111.111 remote-as 100 | Configure neighbor in remote-as 100 |
| (config-router)#neighbor 111.111.111.111 update-source lo | Configure neighbor with update-source loopback |
| (config-router)#address-family vpnv4 unicast | Enter VPNv4 Address family mode |
| (config-router-af)#neighbor 111.111.111.111 activate | mode Activate VPNv4 neighbor |
| (config-router-af)#exit-address-family | Exit from Address Family configuration |
| (config-router)#address-family l2vpn vpls | Enter VPLS Address family mode |
| (config-router-af)#neighbor 111.111.111.111 activate | Activate VPLS neighbor |
| (config-router-af)#exit-address-family | Exit from Address Family configuration |
| (config-router)#address-family vrf l3vpnvr300 | Configure VRF address family |
| (config-router-af)#redistribute connected | Redistribute connected addresses |
| (config-router-af)#exit-address-family | Exit from Address Family configuration |
| (config-router)#exit | Exit from router mode |
| (config)#mpls vplsldp100 100 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig)# vpls-peer 111.111.111.111 | Configuring VPLS mesh peers |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |

| | |
|---|--|
| (config)#mpls l2-circuit VPWS400 400 111.111.111.111 | Configuring VPWS instance with name and VPWS ID |
| (config)#mpls vpls vplsbgp200 200 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls)#signaling bgp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig)#ve-id 201 | Configure VE ID, which is mandatory for BGP VPLS, otherwise, signaling does not take place. VE ID should be unique per VPLS instance |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |
| (config-if)#interface xe12.100 switchport | Enter sub interface mode |
| (config-if)#split-horizon group access1 | Configure split-horizon group on sub-interface |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls vplsldp100 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode |
| (config-if)#interface xe12.400 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 400 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpws | Access VPWS under sub interface |
| (config-acc-if-vpws)#mpls-vpws VPWS400 | Associating the VPWS Instance to the attachment circuit interface. |
| (config-acc-if-vpws)#exit | Exit from access mode |
| (config-if)#interface xe12.200 switchport | Enter sub interface mode |
| (config-if)#split-horizon group access1 | Configure split-horizon group on sub-interface |
| (config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls vplsbgp200 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode |
| (config-if)#interface xe12.300 | Enter sub interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#ip vrf forwarding l3vpnvr300 | Attaching xe12.300 to as part of l3vpnvr300 |
| (config-if)#ip address 210.210.210.1/24 | Configure the IP address of the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the configuration |
| (config)#end | Return to privilege mode |

Validation

PE1

```

---
---
PE1#show ip ospf neighbor

Total number of full neighbors: 4
OSPF process 100 VRF(default):
Neighbor ID      Pri   State           Dead Time   Address      Interface     Instance ID
11.11.11.11      1     Full/Backup     00:00:30   10.0.0.20   po1000.1000   0
11.11.11.11      1     Full/Backup     00:00:34   10.0.1.20   xe13          0
33.33.33.33      1     Full/Backup     00:00:34   16.0.0.20   xe4           0
33.33.33.33      1     Full/Backup     00:00:34   16.0.1.20   xe4.1001      0
PE1#

PE1#show ip ospf interface brief
Interface      PID   Area           Intf ID     Cost  State         Neighbors   Status
lo             100   0.0.0.0        1           1     Loopback     0           Up

Interface      PID   Area           Intf ID     Cost  State         Neighbors   Status
xe4            100   0.0.0.0        10005       10    DR           1           Up

Interface      PID   Area           Intf ID     Cost  State         Neighbors   Status
xe13           100   0.0.0.0        10014       10    DR           1           Up

Interface      PID   Area           Intf ID     Cost  State         Neighbors   Status
xe4.1001       100   0.0.0.0        327844841   10    DR           1           Up

Interface      PID   Area           Intf ID     Cost  State         Neighbors   Status
po1000.1000    100   0.0.0.0        524289000   10    DR           1           Up
PE1#

PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name      My Role      State          KeepAlive  UpTime
11.11.11.11           xe13         Active       OPERATIONAL   30          00:00:33
33.33.33.33           xe4.1001     Active       OPERATIONAL   30          00:00:33
222.222.222.222       xe4.1001     Passive      OPERATIONAL   30          00:00:22

PE1#show bgp neighbors
BGP neighbor is 222.222.222.222, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 111.111.111.111, remote router ID 222.222.222.222
  BGP state = Established, up for 00:00:27
  Last read 00:00:08, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family L2VPN VPLS: advertised and received
    Address family L2VPN EVPN: advertised and received
    Address family IPv6 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Address family IPv6 Labeled Unicast: advertised and received
  Received 11 messages, 0 notifications, 0 in queue
  Sent 13 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo

```

```

For address family: VPNv4 Unicast
BGP table version 2, neighbor version 2
Index 1, Offset 0, Mask 0x2
AIGP is enabled
Community attribute sent to this neighbor (both)
Large Community attribute sent to this neighbor
1 accepted prefixes
1 announced prefixes

```

```

For address family: L2VPN VPLS
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (both)
Large Community attribute sent to this neighbor
0 accepted prefixes
1 announced prefixes

```

```

Connections established 1; dropped 0
Local host: 111.111.111.111, Local port: 179
Foreign host: 222.222.222.222, Foreign port: 35033
Nextthop: 111.111.111.111
Nextthop global: ::
Nextthop local: ::
BGP connection: non shared network

```

```
PE1#
```

```
PE1#show mpls vpls
```

| Name | VPLS-ID | Type | MPeers | SPeers | SIG- |
|-------------------|---------|----------|--------|--------|------|
| Protocol Learning | | | | | |
| vplsldp100 | 100 | Ethernet | 1 | 0 | LDP |
| nabled | | | | | |
| vplsbgp200 | 200 | Ethernet | 1 | 0 | BGP |
| nabled | | | | | |

```
PE1#
```

```
PE1#show mpls vpls mesh
```

```

(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 100 | 222.222.222.222 | 25611 | 26253 | xe13 | 26246 | 2/Up | 1 L |
| DP | Active | 00:56:04 | - | | | | |
| 200 | 222.222.222.222 | 25611 | 25608 | xe13 | 25607 | 2/Up | 3 BGP |
| | Active | 00:56:10 | - | | | | |

```
PE1#show mpls vc-table
```

```

(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VC-ID | Access-Intf | Network-Intf | In label | Out Label | Tunnel- |
|-------|-----------------|--------------|----------|-----------|---------|
| Label | Nextthop | Status | UpTime | | |
| 400 | xe8.400 | - | 26290 | 26252 | N/A |
| (e) | 222.222.222.222 | Active | - | | |

```
PE1#show mpls l2-circuit
```

```

MPLS Layer-2 Virtual Circuit: VPWS400, id: 400 PW-INDEX: 2 service-tpid: dot1q
Endpoint: 222.222.222.222
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe8.400
Subinterface Match Criteria(s) :
dot1q 400

```

```

Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active

```

```
PE1#
```

```
PE1#show mpls vrf-forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, B - BGP FTN
(m) - Service mapped over multipath transport

```

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|------|------------------|--------|-----------|-----|-------------|-----------|------|
| Intf | Nexthop | | | | | | |
| B> | 210.210.210.0/24 | 1 | - | - | LSP_DEFAULT | 25664 | - |
| | 222.222.222.222 | | | | | | |
| B> | 210::/64 | 2 | - | - | LSP_DEFAULT | 25664 | - |
| | 222.222.222.222 | | | | | | |

```
PE1#
```

```
PE1#show mpls forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

| Code | FEC | FTN-ID | Nhife-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|-------------------------|----------------|--------|-----------------|-----------|-----|----------|-----------|------|
| Intf | ELC Nexthop | | | | | | | |
| L> | 9.0.0.0/24 | 1 | 73 (e) | | | | | |
| | | | 52 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe13 | No | 10.0.1.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | (via |
| | | | | | | | | |
| 16.0.0.20 ,label 25612) | | | | | | | | |
| | | | 53 | - | Yes | LSP_ | | |
| DEFAULT | 3 po1000.1000 | No | 10.0.0.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | (via |
| | | | | | | | | |
| 16.0.0.20 ,label 25612) | | | | | | | | |
| L> | 11.0.0.0/24 | 2 | 73 (e) | | | | | |
| | | | 52 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe13 | No | 10.0.1.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | (via |
| | | | | | | | | |
| 16.0.0.20 ,label 25612) | | | | | | | | |
| | | | 53 | - | Yes | LSP_ | | |
| DEFAULT | 3 po1000.1000 | No | 10.0.0.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | (via |
| | | | | | | | | |
| 16.0.0.20 ,label 25612) | | | | | | | | |
| L> | 11.11.11.11/32 | 3 | 100 (e) | | | | | |
| | | | 52 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe13 | No | 10.0.1.20 | | | | | |
| | | | 53 | - | No | LSP_ | | |
| DEFAULT | 3 po1000.1000 | No | 10.0.0.20 | | | | | |
| | | | 53 | - | Yes | LSP_ | | |
| DEFAULT | 3 po1000.1000 | No | 10.0.0.20 | | | | | |
| | | | 52 | - | No | LSP_ | | |
| DEFAULT | 3 xe13 | No | 10.0.1.20 | | | | | |
| L> | 15.0.0.0/24 | 4 | 77 (e) | | | | | |
| | | | 56 | - | Yes | LSP_ | | |
| DEFAULT | 25600 xe13 | No | 10.0.1.20 | | | | | |
| | | | 71 | - | No | LSP_ | | |
| DEFAULT | 0 xe4 | No | 222.222.222.222 | | | | | |

```

16.0.0.20 ,label 25612)                                     (via
                    57      -      Yes  LSP_
DEFAULT 25600      po1000.1000 No  10.0.0.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No  222.222.222.222

16.0.0.20 ,label 25612)                                     (via
L> 15.0.1.0/24      5          80(e)
                    60      -      Yes  LSP_
DEFAULT 25601      xe13        No  10.0.1.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No  222.222.222.222

16.0.0.20 ,label 25612)                                     (via
                    61      -      Yes  LSP_
DEFAULT 25601      po1000.1000 No  10.0.0.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No  222.222.222.222

16.0.0.20 ,label 25612)                                     (via
L> 17.0.0.0/24      6          84(e)
                    16      -      Yes  LSP_
DEFAULT 3          xe4          No  16.0.0.20
                    82      -      No   LSP_
DEFAULT 26249      po1000.1000 No  222.222.222.222

10.0.0.20 ,label 25611)                                     (via
                    17      -      Yes  LSP_
DEFAULT 3          xe4.1001    No  16.0.1.20
                    82      -      No   LSP_
DEFAULT 26249      po1000.1000 No  222.222.222.222

10.0.0.20 ,label 25611)                                     (via
L> 18.0.0.0/24      7          88(e)
                    19      -      Yes  LSP_
DEFAULT 25608      xe4          No  16.0.0.20
                    86      -      No   LSP_
DEFAULT 0          po1000.1000 No  222.222.222.222

10.0.0.20 ,label 25611)                                     (via
                    20      -      Yes  LSP_
DEFAULT 25608      xe4.1001    No  16.0.1.20
                    86      -      No   LSP_
DEFAULT 0          po1000.1000 No  222.222.222.222

10.0.0.20 ,label 25611)                                     (via
L> 18.0.1.0/24      8          91(e)
                    22      -      Yes  LSP_
DEFAULT 25609      xe4          No  16.0.0.20
                    86      -      No   LSP_
DEFAULT 0          po1000.1000 No  222.222.222.222

10.0.0.20 ,label 25611)                                     (via
                    23      -      Yes  LSP_
DEFAULT 25609      xe4.1001    No  16.0.1.20
                    86      -      No   LSP_
DEFAULT 0          po1000.1000 No  222.222.222.222

10.0.0.20 ,label 25611)                                     (via
L> 22.22.22.22/32   9          94(e)
                    64      -      Yes  LSP_
DEFAULT 25607      xe13        No  10.0.1.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No  222.222.222.222

16.0.0.20 ,label 25612)                                     (via
                    65      -      Yes  LSP_
DEFAULT 25607      po1000.1000 No  10.0.0.20
                    71      -      No   LSP_
DEFAULT 0          xe4          No  222.222.222.222

```


(via

```
16.0.0.20 ,label 25612)
L> 33.33.33.33/32 10 102(e)
16 - Yes LSP_
DEFAULT 3 xe4 No 16.0.0.20
17 - No LSP_
DEFAULT 3 xe4.1001 No 16.0.1.20
17 - Yes LSP_
DEFAULT 3 xe4.1001 No 16.0.1.20
16 - No LSP_
DEFAULT 3 xe4 No 16.0.0.20
L> 44.44.44.44/32 11 98(e)
28 - Yes LSP_
DEFAULT 25611 xe4 No 16.0.0.20
96 - No LSP_
DEFAULT 26251 po1000.1000 No 222.222.222.222

(via
10.0.0.20 ,label 25611)
29 - Yes LSP_
DEFAULT 25611 xe4.1001 No 16.0.1.20
96 - No LSP_
DEFAULT 26251 po1000.1000 No 222.222.222.222

(via
10.0.0.20 ,label 25611)
L> 222.222.222.222/32 12 104(e)
68 - Yes LSP_
DEFAULT 25611 xe13 Yes 10.0.1.20
31 - No LSP_
DEFAULT 25612 xe4 Yes 16.0.0.20
69 - Yes LSP_
DEFAULT 25611 po1000.1000 Yes 10.0.0.20
31 - No LSP_
DEFAULT 25612 xe4 Yes 16.0.0.20
31 - Yes LSP_
DEFAULT 25612 xe4 Yes 16.0.0.20
69 - No LSP_
DEFAULT 25611 po1000.1000 Yes 10.0.0.20
32 - Yes LSP_
DEFAULT 25612 xe4.1001 Yes 16.0.1.20
69 - No LSP_
DEFAULT 25611 po1000.1000 Yes 10.0.0.20
B> 220::/64 13 49 - LSP_DEFAULT 25665 -
No 222.222.222.222

PE1#
```

```
PE1#show interface counters rate mbps
```

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-------------|---------|--------|---------|--------|
| po1000 | 350.55 | 4850 | 250.40 | 3466 |
| po1000.1000 | 350.55 | 4851 | 250.39 | 3465 |
| xe0 | 0.00 | 0 | 0.00 | 0 |
| xe2 | 0.00 | 0 | 0.00 | 0 |
| xe4 | 400.63 | 5546 | 500.68 | 6932 |
| xe4.1001 | 400.63 | 5545 | 250.40 | 3465 |
| xe8 | 997.81 | 13858 | 997.81 | 13858 |
| xe8.400 | 997.80 | 13858 | 997.80 | 13858 |
| xe12 | 350.56 | 4850 | 250.40 | 3466 |
| xe13 | 250.28 | 3467 | 250.29 | 3466 |

```
PE1#
```

P1

```
---

P1#show ip ospf interface brief
Interface  PID  Area          Intf ID  Cost  State          Neighbors  Status
```

```

lo          100  0.0.0.0      1          1      Loopback      0          Up

Interface   PID   Area          Intf ID    Cost   State          Neighbors   Status
ge8          100  0.0.0.0      10009     20    Backup         1           Up

Interface   PID   Area          Intf ID    Cost   State          Neighbors   Status
ge13         100  0.0.0.0      10014     10    Backup         1           Up

Interface   PID   Area          Intf ID    Cost   State          Neighbors   Status
po1000.1000  100  0.0.0.0      524289000 10    Backup         1           Up

Interface   PID   Area          Intf ID    Cost   State          Neighbors   Status
po2000.100   100  0.0.0.0      557056100 10    Backup         1           Up

P1#

P1#show ip ospf neighbor

Total number of full neighbors: 4
OSPF process 100 VRF(default):
Neighbor ID   Pri   State          Dead Time   Address        Interface      Instance ID
222.222.222.222 1    Full/DR        00:00:38    9.0.0.20       ge8             0
111.111.111.111 1    Full/DR        00:00:36    10.0.0.10      po1000.1000    0
111.111.111.111 1    Full/DR        00:00:34    10.0.1.10      ge13            0
22.22.22.22    1    Full/DR        00:00:38    11.0.0.20      po2000.100     0
P1#

P2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name      My Role      State          KeepAlive  UpTime
222.222.222.222      ge6.1001     Passive      OPERATIONAL   30           00:00:21
11.11.11.11          po2000.100   Active       OPERATIONAL   30           00:01:27
33.33.33.33          po2000.100   Passive      OPERATIONAL   30           00:00:14

P1#show interface counters rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
ge8          190.14   2634    250.19   3464
ge10         0.00     0        0.00     0
ge12         250.30   3465    350.42   4848
ge13         250.19   3465    250.19   3465
po1000       250.30   3465    350.42   4849
po1000.1000  250.31   3465    350.42   4848
po2000       410.50   5680    250.30   3464
po2000.100   410.50   5681    250.30   3464
xe23         220.26   3048    120.15   1663
xe24         190.23   2632    130.16   1801
xe26         0.00     0        0.00     0
P1#

P1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

LDP ilm-ecmp - enabled
Code  FEC/VRF/L2CKT  ILM-ID      In-Label    Out-Label    In-Intf      Out-
Intf/VRF      Nexthop     pri         LSP-Type     UpTime       UPStr-peers

L> 111.111.111.111/32

```

| | | | | | | | |
|---------|------------------------|-------------|----------|-------|-----------|-------------|------------|
| 22 | 25621 | 0 | N/A | ge13 | 10.0.1.10 | Yes | LSP_ |
| DEFAULT | 00:04:23 | 1 | | | | | |
| | No | LSP_DEFAULT | 25621 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | | | 00:04:23 | | | | |
| | Yes | LSP_DEFAULT | 25621 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | | | 00:04:23 | | | | |
| | No | LSP_DEFAULT | 25621 | 0 | N/A | ge13 | 10.0.1.10 |
| | | | 00:04:23 | | | | |
| L> | 16.0.1.0/24 | 9 | 25608 | 0 | N/A | ge13 | 10.0.1.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25608 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 15.0.0.0/24 | 3 | 25602 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25602 | 0 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 15.0.1.0/24 | 2 | 25601 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25601 | 0 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 15.0.0.0/24 | 1 | 25600 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25600 | 0 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | 1 | | | |
| L> | 16.0.0.0/24 | 5 | 25604 | 0 | N/A | ge13 | 10.0.1.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25604 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 15.0.1.0/24 | 4 | 25603 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25603 | 0 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 18.0.1.0/24 | 7 | 25606 | 25604 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25606 | 0 | N/A | ge8 | 9.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25606 | 3 | N/A | ge8 | 44.44.44.4 |
| 4 | No | LSP_DEFAULT | 00:04:23 | | | | |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 18.0.0.0/24 | 6 | 25605 | 25603 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25605 | 0 | N/A | ge8 | 9.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25605 | 3 | N/A | ge8 | 44.44.44.4 |
| 4 | No | LSP_DEFAULT | 00:04:23 | | | | |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 22.22.22.22/32 | 8 | 25607 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25607 | 26267 | N/A | ge8 | 9.0.0.20 |
| | No | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 111.111.111.111/32 | | | | | | |
| 17 | 25616 | 0 | N/A | ge13 | 10.0.1.10 | Yes | LSP_ |
| DEFAULT | 00:04:23 | | | | | | |
| | No | LSP_DEFAULT | 25616 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | | | 00:04:23 | | | | |
| | Yes | LSP_DEFAULT | 25616 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | | | 00:04:23 | | | | |
| | No | LSP_DEFAULT | 25616 | 0 | N/A | ge13 | 10.0.1.10 |
| | | | 00:04:23 | | | | |
| L> | 16.0.1.0/24 | 14 | 25613 | 0 | N/A | ge13 | 10.0.1.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25613 | 0 | N/A | po1000.1000 | 10.0.0.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| L> | 18.0.0.0/24 | 11 | 25610 | 25603 | N/A | po2000.100 | 11.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |
| | | | 25610 | 0 | N/A | ge8 | 9.0.0.20 |
| | Yes | LSP_DEFAULT | 00:04:23 | 1 | | | |
| L> | 17.0.0.0/24 | 10 | 25609 | 26245 | N/A | ge13 | 10.0.1.10 |
| | Yes | LSP_DEFAULT | 00:04:23 | | | | |

| | | | | | | | |
|---------|------------------------|-----------------|----------|------------|-----------|-------------|-------------|
| | | | 25609 | 26245 | N/A | po1000.1000 | 10.0.0.10 |
| L> | 16.0.0.0/24 | Yes LSP_DEFAULT | 00:04:23 | 25612 | 0 | N/A | ge13 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25612 | 0 | N/A | po1000.1000 |
| L> | 222.222.222.222/32 | Yes LSP_DEFAULT | 00:04:23 | | | | |
| 12 | 25611 | 25607 | N/A | po2000.100 | 11.0.0.20 | Yes | LSP_ |
| DEFAULT | 00:04:23 | | | | | | |
| | | Yes LSP_DEFAULT | 00:04:23 | 25611 | 0 | N/A | ge8 |
| 4 | | No LSP_DEFAULT | 00:04:23 | 25611 | 25624 | N/A | ge8 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 17.0.0.0/24 | 15 | 25614 | 26245 | N/A | ge13 | 10.0.1.10 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25614 | 3 | N/A | ge8 |
| 4 | | No LSP_DEFAULT | 00:04:23 | | | | 44.44.44.4 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| | | Yes LSP_DEFAULT | 00:04:23 | 25614 | 26245 | N/A | po1000.1000 |
| 4 | | No LSP_DEFAULT | 00:04:23 | 25614 | 3 | N/A | ge8 |
| | | | | | | | 44.44.44.4 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 33.33.33.33/32 | 16 | 25615 | 26249 | N/A | ge13 | 10.0.1.10 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25615 | 25622 | N/A | ge8 |
| 4 | | No LSP_DEFAULT | 00:04:23 | | | | 44.44.44.4 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| | | Yes LSP_DEFAULT | 00:04:23 | 25615 | 26249 | N/A | po1000.1000 |
| 4 | | No LSP_DEFAULT | 00:04:23 | 25615 | 25622 | N/A | ge8 |
| | | | | | | | 44.44.44.4 |
| | | | | | | | (via |
| | 9.0.0.20 ,label 26269) | | | | | | |
| L> | 22.22.22.22/32 | 19 | 25618 | 3 | N/A | po2000.100 | 11.0.0.20 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25618 | 26267 | N/A | ge8 |
| | | No LSP_DEFAULT | 00:04:23 | 25617 | 25604 | N/A | po2000.100 |
| L> | 18.0.1.0/24 | 18 | 25617 | 0 | N/A | ge8 | 11.0.0.20 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25617 | 0 | N/A | ge8 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25619 | 26249 | N/A | po1000.1000 |
| L> | 33.33.33.33/32 | 20 | 25619 | 26249 | N/A | ge13 | 10.0.1.10 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25619 | 26249 | N/A | po1000.1000 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 25606 | N/A | po2000.100 |
| L> | 44.44.44.44/32 | 21 | 25620 | 26269 | N/A | ge8 | 11.0.0.20 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 26250 | N/A | ge13 |
| | | No LSP_DEFAULT | 00:04:23 | 25620 | 26269 | N/A | ge8 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 26250 | N/A | po1000.1000 |
| | | No LSP_DEFAULT | 00:04:23 | 25620 | 26269 | N/A | ge8 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 26269 | N/A | ge8 |
| | | No LSP_DEFAULT | 00:04:23 | 25620 | 26250 | N/A | po1000.1000 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25620 | 1 | N/A | po1000.1000 |
| L> | 33.33.33.33/32 | 32 | 25631 | 26249 | N/A | ge13 | 10.0.1.10 |
| | | Yes LSP_DEFAULT | 00:04:23 | 25631 | 26249 | N/A | po1000.1000 |
| | | Yes LSP_DEFAULT | 00:04:23 | | | | |

```

L> 16.0.0.0/24      26      25625      0      N/A      ge13      10.0.1.10
      Yes      LSP_DEFAULT      00:04:23      25625      0      N/A      po1000.1000      10.0.0.10
      Yes      LSP_DEFAULT      00:04:23
L> 15.0.0.0/24      24      25623      3      N/A      po2000.100      11.0.0.20
      Yes      LSP_DEFAULT      00:04:23      25623      0      N/A      ge8      9.0.0.20
      No      LSP_DEFAULT      00:04:23
L> 222.222.222.222/32
23      25622      25607      N/A      po2000.100      11.0.0.20      Yes      LSP_
DEFAULT      00:04:23
      Yes      LSP_DEFAULT      00:04:23      25622      0      N/A      ge8      9.0.0.20
L> 15.0.1.0/24      25      25624      3      N/A      po2000.100      11.0.0.20
      Yes      LSP_DEFAULT      00:04:23      25624      0      N/A      ge8      9.0.0.20
      No      LSP_DEFAULT      00:04:23

```

P2

```

---
P2#show ip ospf interface brief
Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
lo             100      0.0.0.0      1             1      Loopback      0              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
ge6.300        100      0.0.0.0      328171820    10      Backup      1              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
ge6.1001       100      0.0.0.0      328172521    10      Backup      1              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
po2000.100     100      0.0.0.0      557056100    10      DR          1              Up

P2#

P2#show ip ospf neighbor

Total number of full neighbors: 3
OSPF process 100 VRF(default):
Neighbor ID      Pri      State      Dead Time      Address      Interface      Instance ID
11.11.11.11      1      Full/Backup  00:00:32      11.0.0.10    po2000.100     0
222.222.222.222  1      Full/DR     00:00:36      15.0.0.20    ge6.300        0
222.222.222.222  1      Full/DR     00:00:35      15.0.1.20    ge6.1001       0
P2#

P2#show ldp session
Codes: m - MD5 password is not set/unset.
      g - GR configuration not set/unset.
      t - TCP MSS not set/unset.
      Session has to be cleared manually

Code  Peer IP Address      IF Name      My Role      State      KeepAlive      UpTime
222.222.222.222      ge6.1001      Passive      OPERATIONAL  30          00:00:21
11.11.11.11          po2000.100    Active      OPERATIONAL  30          00:01:27
33.33.33.33          po2000.100    Passive      OPERATIONAL  30          00:00:14      00:00:14
---
P3#show ip ospf interface brief
Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
lo             100      0.0.0.0      1             1      Loopback      0              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
ge5            100      0.0.0.0      10014         10      Backup      1              Up

Interface      PID      Area      Intf ID      Cost      State      Neighbors      Status
ge5.1001       100      0.0.0.0      328139753     10      Backup      1              Up

```

| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
|------------|-----|---------|------------|------|--------|-----------|--------|
| sa4000.200 | 100 | 0.0.0.0 | 1179648200 | 10 | Backup | 1 | Up |

P3#

P3#show ip ospf neighbor

Total number of full neighbors: 3

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-----------------|-----|---------|-----------|-----------|------------|-------------|
| 111.111.111.111 | 1 | Full/DR | 00:00:38 | 16.0.0.10 | ge5 | 0 |
| 111.111.111.111 | 1 | Full/DR | 00:00:32 | 16.0.1.10 | ge5.1001 | 0 |
| 44.44.44.44 | 1 | Full/DR | 00:00:36 | 17.0.0.10 | sa4000.200 | 0 |

P3#

P3#show ldp session

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|------------|---------|-------------|-----------|----------|
| 44.44.44.44 | sa4000.200 | Passive | OPERATIONAL | 30 | 00:02:00 |
| 111.111.111.111 | ge5 | Passive | OPERATIONAL | 30 | 00:00:32 |
| 22.22.22.22 | ge5 | Active | OPERATIONAL | 30 | 00:00:14 |
| 11.11.11.11 | ge5 | Active | OPERATIONAL | 30 | 00:00:20 |
| 222.222.222.222 | sa4000.200 | Passive | OPERATIONAL | 30 | 00:00:13 |

P3#

P4

P4#show ip ospf interface brief

| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
|-------------|-----|---------|------------|------|----------|-----------|--------|
| lo | 100 | 0.0.0.0 | 1 | 1 | Loopback | 0 | Up |
| sa3000 | 100 | 0.0.0.0 | 203000 | 10 | Backup | 1 | Up |
| sa3000.1001 | 100 | 0.0.0.0 | 1146881001 | 10 | Backup | 1 | Up |
| sa4000.200 | 100 | 0.0.0.0 | 1179648200 | 10 | DR | 1 | Up |

P4#

P4#show ip ospf neighbor

Total number of full neighbors: 3

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-----------------|-----|-------------|-----------|-----------|-------------|-------------|
| 33.33.33.33 | 1 | Full/Backup | 00:00:34 | 17.0.0.20 | sa4000.200 | 0 |
| 222.222.222.222 | 1 | Full/DR | 00:00:36 | 18.0.0.20 | sa3000 | 0 |
| 222.222.222.222 | 1 | Full/DR | 00:00:34 | 18.0.1.20 | sa3000.1001 | 0 |

P4#

P4#show ldp session

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|------------|---------|-------------|-----------|----------|
| 222.222.222.222 | sa3000 | Passive | OPERATIONAL | 30 | 00:00:40 |
| 33.33.33.33 | sa4000.200 | Active | OPERATIONAL | 30 | 00:02:00 |
| 11.11.11.11 | sa4000.200 | Active | OPERATIONAL | 30 | 00:00:14 |

P4#

PE2

PE2#show ip ospf neighbor

Total number of full neighbors: 5

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------|-----------|---------|-----------|-------------|
|-------------|-----|-------|-----------|---------|-----------|-------------|

```

11.11.11.11      1    Full/Backup    00:00:36    9.0.0.10    ge9          0
22.22.22.22      1    Full/Backup    00:00:31    15.0.0.10   xe14.1002    0
22.22.22.22      1    Full/Backup    00:00:34    15.0.1.10   xe14.1001    0
44.44.44.44      1    Full/Backup    00:00:38    18.0.0.10   sa3000       0
44.44.44.44      1    Full/Backup    00:00:31    18.0.1.10   sa3000.1001  0
PE2#

```

PE2#show ip ospf interface brief

| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
|-------------|-----|---------|------------|------|----------|-----------|--------|
| lo | 100 | 0.0.0.0 | 1 | 1 | Loopback | 0 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| ge9 | 100 | 0.0.0.0 | 10018 | 20 | DR | 1 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| sa3000 | 100 | 0.0.0.0 | 203000 | 10 | DR | 1 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| xe14.1001 | 100 | 0.0.0.0 | 328434665 | 10 | DR | 1 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| xe14.1002 | 100 | 0.0.0.0 | 328434666 | 10 | DR | 1 | Up |
| Interface | PID | Area | Intf ID | Cost | State | Neighbors | Status |
| sa3000.1001 | 100 | 0.0.0.0 | 1146881001 | 10 | DR | 1 | Up |

PE2#

PE2#show ldp session

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|-----------|---------|-------------|-----------|----------|
| 44.44.44.44 | sa3000 | Active | OPERATIONAL | 30 | 00:00:41 |
| 11.11.11.11 | ge9 | Active | OPERATIONAL | 30 | 00:00:18 |
| 22.22.22.22 | xe14.1002 | Active | OPERATIONAL | 30 | 00:00:25 |
| 111.111.111.111 | xe14.1001 | Active | OPERATIONAL | 30 | 00:00:22 |
| 33.33.33.33 | sa3000 | Active | OPERATIONAL | 30 | 00:00:14 |

PE2#

PE2#show bgp neighbors

```

BGP neighbor is 111.111.111.111, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 222.222.222.222, remote router ID 111.111.111.111
  BGP state = Established, up for 00:00:27
  Last read 00:00:01, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family VPNv4 Unicast: advertised and received
    Address family L2VPN VPLS: advertised and received
    Address family L2VPN EVPN: advertised and received
    Address family IPv6 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
    Address family IPv6 Labeled Unicast: advertised and received
  Received 12 messages, 0 notifications, 0 in queue
  Sent 11 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo

```

For address family: VPNv4 Unicast

```

  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  1 announced prefixes

```

```

For address family: L2VPN VPLS
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (both)
Large Community attribute sent to this neighbor
0 accepted prefixes
1 announced prefixes

```

```

Connections established 1; dropped 0
Local host: 222.222.222.222, Local port: 35033
Foreign host: 111.111.111.111, Foreign port: 179
Nextthop: 222.222.222.222
Nextthop global: ::
Nextthop local: ::
BGP connection: non shared network

```

PE2#

PE2#show mpls vpls

| Name | VPLS-ID | Type | MPeers | SPeers | SIG- | |
|-------------------|---------|----------|--------|--------|------|---|
| Protocol Learning | | | | | | |
| vplsldp100 | 100 | Ethernet | 1 | 0 | LDP | E |
| nabled | | | | | | |
| vplsbgp200 | 200 | Ethernet | 1 | 0 | BGP | E |
| nabled | | | | | | |

PE2#

PE2#show mpls vpls mesh

(m) - Service mapped over multipath transport

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | |
|--------------|-----------------|--------------|----------|--------------|-----------|---------|----------|---|
| SIG-Protocol | Status | | | | | | | |
| 100 | 111.111.111.111 | 25614 | 26246 | sa3000.1001 | 26253 | 2/Up | 1 | L |
| DP | Active | | | | | | | |
| 200 | 111.111.111.111 | 25614 | 25607 | sa3000.1001 | 25608 | 2/Up | 3 | B |
| GP | Active | | | | | | | |

PE2#

PE2#show mpls vc-table

(m) - Service mapped over multipath transport

(e) - Service mapped over LDP ECMP

| VC-ID | Vlan-ID | Inner-Vlan-ID | Access-Intf | Network-Intf | Out Label | Tunnel- |
|-------|-----------------|---------------|-------------|--------------|-----------|---------|
| Label | Nextthop | Status | | | | |
| 400 | N/A | N/A | xe12.400 | N/A (e) | 26252 | N/A |
| (e) | 111.111.111.111 | Active | | | | |

PE2#

PE2#show mpls l2-circuit

MPLS Layer-2 Virtual Circuit: VPWS400, id: 400 PW-INDEX: 2 service-tpid: dot1q

Endpoint: 111.111.111.111

Control Word: 0

Flow Label Status: Disabled, Direction: None, Static: No

MPLS Layer-2 Virtual Circuit Group: none

Bound to interface: xe12.400

Subinterface Match Criteria(s) :

dot1q 400

Virtual Circuit Type: Ethernet VLAN

Virtual Circuit is configured as Primary

Virtual Circuit is configured as Active

Virtual Circuit is active

PE2#

PE2#show mpls vrf-forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, B - BGP FTN

(m) - Service mapped over multipath transport

| Code Intf | FEC Nexthop | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|--------------|-------------------------------------|--------|-----------|-----|-------------|-----------|------|
| B> | 110.110.110.0/24 111.111.111.111 | 2 | - | - | LSP_DEFAULT | 25664 | - |
| B> | 110::/64 111.111.111.111 | 1 | - | - | LSP_DEFAULT | 25664 | - |

PE2#

PE2#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

| Code Intf | FEC ELC Nexthop | FTN-ID UpTime | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|------------------------|--------------------|------------------|-----------------|-----------|----------|----------|-----------|------|
| L> | 10.0.0.0/24 | 8 | 56(e) | | | | | |
| DEFAULT | 25617 xe14.1001 | No | 54 | - | Yes | LSP_ | | |
| | | | 15.0.1.10 | | 00:03:19 | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 25617 xe14.1002 | No | 55 | - | Yes | LSP_ | | |
| | | | 15.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | Yes | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 0 ge9 | No | 38 | - | No | LSP_ | | |
| | | | 111.111.111.111 | | | | | (via |
| 9.0.0.10 ,label 25642) | | | | | | | | |
| L> | 10.0.1.0/24 | 9 | 59(e) | | | | | |
| DEFAULT | 25618 xe14.1001 | No | 57 | - | Yes | LSP_ | | |
| | | | 15.0.1.10 | | - | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 25618 xe14.1002 | No | 58 | - | Yes | LSP_ | | |
| | | | 15.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | Yes | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 0 ge9 | No | 38 | - | No | LSP_ | | |
| | | | 111.111.111.111 | | | | | (via |
| 9.0.0.10 ,label 25642) | | | | | | | | |
| L> | 11.0.0.0/24 | 11 | 63(e) | | | | | |
| DEFAULT | 3 xe14.1001 | No | 60 | - | Yes | LSP_ | | |
| | | | 15.0.1.10 | | 00:03:19 | | | |
| DEFAULT | 3 xe14.1002 | No | 61 | - | Yes | LSP_ | | |
| | | | 15.0.0.10 | | | | | |
| DEFAULT | 26255 ge9 | No | 62 | - | No | LSP_ | | |
| | | | 111.111.111.111 | | | | | (via |
| 9.0.0.10 ,label 25642) | | | | | | | | |
| L> | 11.11.11.11/32 | 10 | 68(e) | | | | | |
| DEFAULT | 25619 xe14.1001 | No | 66 | - | Yes | LSP_ | | |
| | | | 15.0.1.10 | | - | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 25619 xe14.1002 | No | 67 | - | Yes | LSP_ | | |
| | | | 15.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | No | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 33 | - | Yes | LSP_ | | |
| | | | 9.0.0.10 | | | | | |
| DEFAULT | 3 ge9 | No | 40 | - | No | LSP_ | | |

```

DEFAULT 26256 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
L> 16.0.0.0/24 1 42(e)
1 - Yes LSP_
DEFAULT 25604 sa3000.1001 No 18.0.1.10 00:03:19
38 - No LSP_
DEFAULT 0 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
2 - Yes LSP_
DEFAULT 25604 sa3000 No 18.0.0.10
38 - No LSP_
DEFAULT 0 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
L> 16.0.1.0/24 2 43(e)
4 - Yes LSP_
DEFAULT 25605 sa3000.1001 No 18.0.1.10 -
38 - No LSP_
DEFAULT 0 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
5 - Yes LSP_
DEFAULT 25605 sa3000 No 18.0.0.10
38 - No LSP_
DEFAULT 0 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
L> 17.0.0.0/24 3 53(e)
7 - Yes LSP_
DEFAULT 3 sa3000.1001 No 18.0.1.10 00:03:19
52 - No LSP_
DEFAULT 3 ge9 No 33.33.33.33
9.0.0.10 ,label 25641)
8 - Yes LSP_
DEFAULT 3 sa3000 No 18.0.0.10
52 - No LSP_
DEFAULT 3 ge9 No 33.33.33.33
9.0.0.10 ,label 25641)
L> 22.22.22.22/32 12 69(e)
60 - Yes LSP_
DEFAULT 3 xe14.1001 No 15.0.1.10 -
61 - No LSP_
DEFAULT 3 xe14.1002 No 15.0.0.10
61 - Yes LSP_
DEFAULT 3 xe14.1002 No 15.0.0.10
60 - No LSP_
DEFAULT 3 xe14.1001 No 15.0.1.10
L> 33.33.33.33/32 4 45(e)
10 - Yes LSP_
DEFAULT 25606 sa3000.1001 No 18.0.1.10 00:03:19
44 - No LSP_
DEFAULT 26263 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
11 - Yes LSP_
DEFAULT 25606 sa3000 No 18.0.0.10
44 - No LSP_
DEFAULT 26263 ge9 No 111.111.111.111
9.0.0.10 ,label 25642)
L> 44.44.44.44/32 5 31(e)
7 - Yes LSP_
DEFAULT 3 sa3000.1001 No 18.0.1.10 -
8 - No LSP_
DEFAULT 3 sa3000 No 18.0.0.10
8 - Yes LSP_
DEFAULT 3 sa3000 No 18.0.0.10

```

```

DEFAULT 3          sa3000.1001 No 7 - No LSP_
L> 111.111.111.111/32 6 73(e) 71 - Yes LSP_
                                36 - 00:03:19
DEFAULT 25622      xe14.1001 Yes 15.0.1.10 No LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 Yes LSP_
                                72 - 15.0.0.10
DEFAULT 25622      xe14.1002 Yes 15.0.0.10 No LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 Yes LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 No LSP_
                                17 - 18.0.0.10
DEFAULT 25614      sa3000 Yes 18.0.0.10 Yes LSP_
                                16 - 18.0.1.10
DEFAULT 25614      sa3000.1001 Yes 18.0.1.10 No LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 Yes LSP_
                                17 - 18.0.0.10
DEFAULT 25614      sa3000 Yes 18.0.0.10 No LSP_
                                36 - 9.0.0.10
DEFAULT 25642      ge9 Yes 9.0.0.10 - LSP_DEFAULT 25665 -
B> 120::/64 7 23 -
    No 111.111.111.111
PE2#

```

PE2#show interface counters rate mbps

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-------------|---------|--------|---------|--------|
| ge4 | 0.00 | 0 | 0.00 | 0 |
| ge5 | 240.28 | 3327 | 210.29 | 2913 |
| ge6 | 0.00 | 0 | 0.00 | 0 |
| ge7 | 260.29 | 3603 | 190.27 | 2635 |
| ge9 | 250.29 | 3466 | 190.22 | 2635 |
| ge10 | 0.00 | 0 | 0.00 | 0 |
| sa3000 | 500.58 | 6930 | 400.56 | 5548 |
| sa3000.1001 | 0.00 | 0 | 200.33 | 2774 |
| xe12 | 997.83 | 13858 | 997.83 | 13858 |
| xe12.400 | 997.83 | 13858 | 997.83 | 13858 |
| xe14 | 250.40 | 3468 | 410.66 | 5685 |
| xe14.1001 | 0.00 | 1 | 58.53 | 811 |
| xe14.1002 | 250.40 | 3467 | 200.33 | 2772 |

PE2#

EVPN Active-Standby

Overview

EVPN Multihoming is a mechanism that allows a host or customer edge (CE) device to be connected to multiple Provider Edge (PE) devices called Multihoming (MH) peers for redundancy and load balancing purposes. This provides high availability and resiliency to the network, ensuring continuous connectivity even in case of a PE device failure.



Note: OcNOS support extends to a maximum of two MH peers.

Multihoming supports two kinds of redundancy, namely 1. All Active 2. Active-Standby.

Till now, OcNOS support All-Active (A-A) only. In OcNOS version 6.4.1, Port-Active mode is supported and in OcNOS version 6.4.2, Single-Active mode is supported in the context of Active-Standby redundancy.

Single-Active

- In this mode, traffic for a specific host or MAC address is handled by only one of the PE devices (MH peers) at a time.
- The other PE devices remain in standby mode, ready to take over if the active PE fails.
- The physical link state (either Physical port or LACP port) on the standby PE remains up, enabling a faster transition to the active role in the event of a failover. The CE devices use different interfaces, including LACP or physical connections, to connect to the Peer MH devices.

Port-Active

- In this mode, traffic for a specific host or MAC address is handled by only one of the PE devices (MH peers) at a time.
- Each MH peer connects through LACP with the same key as the CE devices (similar to A-A redundancy).
- The physical link state (LACP port) on the standby PE is made down, effectively blocking traffic on those ports.
- If a failover occurs, the standby PE must bring up its LACP ports to start forwarding traffic.

IRB Active-Standby: Active-standby mode is also applicable to Integrated Routing and Bridging (IRB) for both L3VPN symmetric and asymmetric modes.

Feature Characteristics

Single-Active standby redundancy mechanisms support both ELAN and ELINE services.

Single-Active ELINE

ELINE refers to Ethernet Line services, where two PEs are cross-connected to each other over an Ethernet link.

In Single-Active ELINE, the primary objective is to achieve redundancy for hosts while also using the same link for data exchange until it fails, at which point it should switch to the secondary or standby link. Here's how it works:

- **MH Host Traffic**

- One of the PE devices (MH peers) acts as the "Active" for the Attachment Circuit (AC) associated with the host. This PE sends and receives traffic to and from the host.
- The other PE acts as the "Standby" for the same AC and does not allow traffic to or from the host.
- The standby PE, despite receiving BUM traffic from the Host device (which is unaware of the cross-connect), blocks this traffic at the standby PE itself, as it operates in a standby role for the AC. Conversely, the active PE allows the flow of traffic.
- **Remote Host Traffic:** Traffic originating from remote hosts destined for the multihomed host is only sent to the active MH peer for the corresponding AC. This ensures that the cross connect is established only with the Active MH peer.

Single-Active ELAN

ELAN stands for Ethernet LAN services, where a group of PEs are interconnected in a multipoint Ethernet network. In Single-Active ELAN, similar to Single-Active ELINE, redundancy for hosts and data exchange over the primary link are priorities, but there are some specific differences for Ethernet LAN (ELAN) scenarios:

- **MH Host Traffic**
 - One of the PE devices (MH peers) is designated as the "Active" for the AC associated with the host. This PE handles sending and receiving traffic to and from remote locations.
 - The other PE acts as the "Standby" role for the same AC. It receives BUM traffic from the host but blocks the traffic. Additionally, it refrains from learning MAC addresses and does not uplift Address Resolution Protocol/Neighbor Discovery (ARP/ND) packets.
 - Unicast traffic from the host will be directed to the active PE, which will then allow the traffic to be sent across the network.
- **Remote Host Unicast Traffic:** Unicast traffic from remote hosts destined for the multihomed host is sent only to the active MH peer for the corresponding AC. This is because the MAC addresses of the host are learned only from the Active MH peer.
- **Remote Host BUM Traffic:** BUM traffic, such as broadcast and multicast packets from the remote PE device, is replicated to both MH PEs. However, only the active PE, which is also designated as a forwarder, allows this traffic to reach the host. The standby PE, classified as a Non-Designated Forwarder, drops the egress traffic.

Port-Active Ethernet LAN (ELAN) and Ethernet LINE (ELINE) are examples of port-active standby redundancy mechanisms.

Port-Active ELINE

Port-Active ELINE enables redundancy and optimized data exchange by designating an active port for traffic handling in multihomed network setups. Here's how it works:

- **Active AC Link:** Among the Multihomed (MH) peers, a designated PE is assigned as "Active" for the AC associated with the host. This PE manages bidirectional traffic to and from the host. In a port-active configuration, all hosts associated with the ESI link remain in the same state, as the Active and Standby status is determined per ESI link.
- **Standby AC Link:** The AC link attached to the host, designated as "Standby," remains operationally down. It serves as a backup link for failover scenarios.
- **MH Host traffic:** BUM and unicast traffic from the host are always directed towards the Active PE because the link towards the Active PE is operational UP. Conversely, the link towards the Standby PE from the host devices is operational DOWN.

- **Remote Host Traffic:** Traffic originating from remote hosts and destined for the multihomed host is directed exclusively to the Active MH node that serves the corresponding AC. This ensures efficient traffic routing and intelligent cross-connection establishment.

Port-Active ELAN

Port-Active ELAN enhances redundancy and efficient data exchange by designating an active port for traffic management in multihomed Ethernet LAN environments. Here's how it works:

- **Active AC Link:** Within the MH peers, one PE is identified as the "Active" entity for the AC. It manages traffic to and from remote locations efficiently.
- **Standby AC Link:** Similar to Port-Active ELINE, the standby AC link attached to the host remains operationally down to ensure effective standby redundancy.
- **MH Host Traffic:** In a port-active scenario, the standby link does not receive any traffic from the host. Only the active link manages incoming traffic from the host. The Active PE also learns and advertises host information to remote locations, including MAC addresses and ARP/ND details.
- **Remote Host Unicast Traffic:** Unicast traffic from remote sources is directed exclusively to the Active MH PE that has advertised the host address, optimizing traffic flow.
- **Remote Host BUM Traffic:** BUM traffic is replicated across all MH nodes. However, egress traffic for BUM packets occurs only from the Active PE. The standby PE drops the traffic since the AC links are operational DOWN.

Benefits

The benefits of Single-Active and Port-Active include enhanced redundancy and fault tolerance for hosts and customer edge devices, efficient data exchange, minimized downtime, and improved network resiliency in multihomed Ethernet Line and Ethernet LAN environments. These mechanisms ensure uninterrupted connectivity and optimized traffic management, contributing to higher availability and improved user experience.

Prerequisites

Here are the prerequisites for configuring EVPN Multihoming:

Ensure EVPN Configuration: Make sure that the EVPN is configured already in the network as it is a requirement for EVPN Multihoming.

Configure Attachment Circuits (AC): Ensure that each CE device is appropriately linked to the PE devices through Attachment Circuits. These circuits must be configured correctly.

Set Up LACP Configuration: To use Link Aggregation Control Protocol (LACP) for multihoming, configure LACP appropriately on the relevant interfaces.

EVPN MPLS Global Configuration: To enable EVPN MPLS features, need to configure global settings, such as enabling EVPN MPLS, defining global VTEP IP addresses, enabling hardware profile filtering for multihoming, and activating EVPN MPLS multihoming functionality. These settings are essential for EVPN and MPLS operation.

Access Port Configuration: Depending on the network's redundancy plan (single-active or port-active), configure access ports, including parameters for load balancing, service carving preferences, and EVPN settings. These configurations are crucial for network access and connectivity in an EVPN environment.

These prerequisites ensure that the network is ready for the implementation of EVPN Multihoming, providing redundancy and load balancing for CE devices.

Configuration

Here are sample configurations for [EVPN MPLS Active-Standby MultiHoming Configuration \(page 1755\)](#) and [EVPN SR Active-Standby Multi-Homing Configuration \(page 1798\)](#), including topology, configuration procedures, and corresponding validations.

For more information on the EVPN MPLS configurations, see the and [EVPN MPLS IRB Configuration](#) chapters in the Multi-Protocol Label Switching Guide.


EVPN MPLS Active-Standby MultiHoming Configuration

This section illustrates the Multi-Homed setup for the EVPN MPLS Active-Standby configuration, showcasing examples for both ELINE and ELAN services with LDP as the underlay MPLS path.

EVPN MPLS Active-Standby MH Topology

The below topology consists of customer edge routers CE1 and CE2, along with IPv4 Provider Edge routers PE1, PE2, PE3, and PE4, all interconnected through the core routers P1 and P2 in the IPv4 MPLS provider network.

Figure 54. EVPN MPLS AS MH Configuration



Customer Edge (CE) Routers Configurations

CE1

The following configuration steps under CE1 are set up to enable VLANs and configure interfaces for carrying VLAN traffic.

| | |
|--|--|
| CE1#configure terminal | Enter configure mode. |
| CE1(config)#bridge 1 protocol ieee vlan-bridge | Set up bridge 1 to use the IEEE VLAN bridge protocol. |
| CE1(config)#vlan 2-100 bridge 1 state enable | Configure VLANs from 2-100 and associate them with bridge 1. |
| CE1(config)#interface xe54 | Enter interface mode xe54. |

| | |
|--|--|
| CE1(config-if)#switchport | Configure the interface xe54 as a Layer 2 switch port. |
| CE1(config-if)#bridge-group 1 | Associate xe54 to bridge 1. |
| CE1(config-if)#switchport mode trunk | Configure xe54 as a trunk port. |
| CE1(config-if)#switchport trunk allowed vlan all | Allow all configured VLANs on the trunk interface xe54. |
| CE1(config-if)#exit | Exit interface mode xe54. |
| CE1(config)#interface po1 | Enter interface mode and configure LAG interface port-channel 1 (po1). |
| CE1(config-if)#switchport | Configures port-channel 1 as a Layer 2 switch port. |
| CE1(config-if)#bridge-group 1 | Associate po1 to bridge 1. |
| CE1(config-if)#switchport mode trunk | Configure po1 as a trunk port. |
| CE1(config-if)#switchport trunk allowed vlan all | Allow all configured VLANs on the trunk port-channel po1. |
| CE1(config-if)#exit | Exit interface mode po1. |
| CE1(config)#interface xe48 | Enter interface mode xe48. |
| CE1(config-if)#lacp timeout short | Configure LACP timeout as short. |
| CE1(config-if)#channel-group 1 mode active | Add member to the LAG interface. |
| CE1(config-if)#exit | Exit interface mode xe48. |
| CE1(config-if)#interface xe50 | Enter interface mode xe50. |
| CE1(config-if)#lacp timeout short | Configure LACP timeout as short. |
| CE1(config-if)#channel-group 1 mode active | Add member to the LAG interface. |
| CE1(config-if)#commit | Commit the transaction. |
| CE1(config-if)#end | Exit interface mode xe50 and configure mode. |

CE2

The following configuration steps under CE2 are set up to enable VLANs and configure interfaces for carrying VLAN traffic.

| | |
|--|--|
| CE2#configure terminal | Enter configure mode. |
| CE2(config)#bridge 1 protocol ieee vlan-bridge | Set up bridge 1 to use the IEEE VLAN bridge protocol. |
| CE2(config)#vlan 2-100 bridge 1 state enable | Configure VLANs from 2-100 and associate them with bridge 1. |
| CE2(config)#interface xe24 | Enter interface mode xe24. |
| CE2(config-if)#switchport | Configure the interface xe24 as a Layer 2 switch port. |
| CE2(config-if)#bridge-group 1 | Associate xe24 to bridge 1. |
| CE2(config-if)#switchport mode trunk | Configure xe24 as a trunk port. |

| | |
|--|--|
| CE2(config-if)#switchport trunk allowed vlan all | Allow all configured VLANs on the trunk interface xe24. |
| CE2(config-if)#exit | Exit interface mode xe24. |
| CE2(config)#interface po1 | Enter interface mode and configure LAG interface port-channel 1 (po1). |
| CE2(config-if)#switchport | Configures port-channel 1 as a Layer 2 switch port. |
| CE2(config-if)#bridge-group 1 | Associate po1 to bridge 1. |
| CE2(config-if)#switchport mode trunk | Configure po1 as a trunk port. |
| CE2(config-if)#switchport trunk allowed vlan all | Allow all configured VLANs on the trunk port-channel po1. |
| CE2(config-if)#exit | Exit interface mode po1. |
| CE2(config)#interface xe22 | Enter interface mode xe22. |
| CE2(config-if)#lacp timeout short | Configure LACP timeout as short. |
| CE2(config-if)#channel-group 1 mode active | Add member to the LAG interface. |
| CE2(config-if)#exit | Exit interface mode xe22. |
| CE2(config-if)#interface xe23 | Enter interface mode xe23. |
| CE2(config-if)#lacp timeout short | Configure LACP timeout as short. |
| CE2(config-if)#channel-group 1 mode active | Add member to the LAG interface. |
| CE2(config-if)#commit | Commit the transaction. |
| CE2(config-if)#end | Exit interface mode xe23 and configure mode. |

PE1 Router Configurations

PE1: Loopback Interface

The configuration on PE1 for a loopback interface with IP address 10.10.10.1/32 secondary is set up to provide IP connectivity for the router.

| | |
|---|---|
| PE1#configure terminal | Enter configure mode. |
| PE1(config)#interface lo | Enter the interface mode for the loopback interface lo. |
| PE1(config-if)#ip address 10.10.10.1/32 secondary | Configure a secondary IP address, 10.10.10.1/32, on the loopback interface. |
| PE1(config-if)#exit | Exit interface mode lo. |
| PE1(config)#commit | Commit the transaction. |

PE1: Global LDP

The configuration on PE1 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|------------------------|----------------------------|
| PE1(config)#router ldp | Enter the Router LDP mode. |
|------------------------|----------------------------|

| | |
|--|---|
| PE1(config-router)#router-id 10.10.10.1 | Set the router ID for LDP to 10.10.10.1. |
| PE1(config-router)#transport-address ipv4 10.10.10.1 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| PE1(config-router)#targeted-peer ipv4 10.10.10.2 | Configure targeted peer for LDP using IPv4 addresses. |
| PE1(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE1(config-router)#targeted-peer ipv4 10.10.10.3 | Configure targeted peer for LDP using IPv4 addresses. |
| PE1(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE1(config-router)#targeted-peer ipv4 10.10.10.4 | Configure targeted peer for LDP using IPv4 addresses. |
| PE1(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE1(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: Global EVPN MPLS Command

The configuration on PE1 for the Global EVPN MPLS, includes activating EVPN MPLS, defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multi-homing, and activating EVPN MPLS multi-homing functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|---|--|
| PE1(config)#evpn mpls enable | Activate the EVPN MPLS functionality on PE1, enabling it to participate in EVPN MPLS services. |
| PE1(config)#commit | Commit candidate configuration to be running configuration. |
| PE1(config)#evpn mpls vtep-ip-global 10.10.10.1 | Configure the global VTEP IP address 10.10.10.1, associating it with the loopback IP. |
| PE1(config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter for EVPN MPLS multi-homing. |
| PE1(config)#evpn mpls multihoming enable | Activate the EVPN MPLS multi-homing functionality, allowing PE1 to support multi-homed EVPN MPLS services. |
| PE1(config)#commit | Commit the transaction. |

PE1: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE1 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|---------------------------------------|--|
| PE1(config)#interface xe1 | Enter interface mode xe1. |
| PE1(config-if)#ip address 10.1.1.1/30 | Configure an IP address, 10.1.1.1/30, on the interface xe1. |
| PE1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE1(config-if)#exit | Exit interface mode xe1. |
| PE1(config)#commit | Commit the transaction. |
| PE1(config)#interface xe2 | Enter interface mode xe2. |
| PE1(config-if)#ip address 10.1.2.1/30 | Configure an IP address, 10.1.2.1/30, on the interface xe2. |
| PE1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE1(config-if)#exit | Exit interface mode xe2. |
| PE1(config)#commit | Commit the transaction. |

PE1: OSPF Configuration

The below configuration is performed to set up OSPF on PE1, specifying the router ID, defining network interfaces, and configuring BFD parameters for efficient routing.

| | |
|--|---|
| PE1(config)#router ospf 100 | Enter the router OSPF mode. Configure PE1 to run OSPF with process ID 100. |
| PE1(config-router)#ospf router-id 10.10.10.1 | Set the OSPF router ID to 10.10.10.1, identifying PE1 within the OSPF network. |
| PE1(config-router)#network 10.10.10.1/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| PE1(config-router)#network 10.1.1.1/30 area 0.0.0.0 | Advertise network address in OSPF. |
| PE1(config-router)#network 10.1.2.1/30 area 0.0.0.0 | Advertise network address in OSPF. |
| PE1(config-router)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval with an interval of 3, a minimum receive interval of 3, and a multiplier of 3. |
| PE1(config-router)#exit | Exit router OSPF mode and return to configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: BGP Configuration

The below BGP configuration on PE1 is established to enable BGP routing with ASN 65010, set the BGP router ID, define iBGP neighbors, configure BFD, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|---|---|
| PE1(config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| PE1(config-router)#bgp router-id 10.10.10.1 | Configure BGP router ID 10.10.10.1, identifying PE1 within the BGP network. |
| PE1(config-router)#neighbor 10.10.10.2 remote-as 65010 | Configure neighbor 10.10.10.2 as an iBGP neighbor with their remote AS number 65010. |
| PE1(config-router)#neighbor 10.10.10.2 update-source lo | Configure neighbor 10.10.10.2 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE1(config-router)#neighbor 10.10.10.3 remote-as 65010 | Configure neighbor 10.10.10.3 as an iBGP neighbor with their remote AS number 65010. |
| PE1(config-router)#neighbor 10.10.10.3 update-source lo | Configure neighbor 10.10.10.3 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE1(config-router)#neighbor 10.10.10.4 remote-as 65010 | Configure neighbor 10.10.10.4 as an iBGP neighbor with their remote AS number 65010. |
| PE1(config-router)#neighbor 10.10.10.4 update-source lo | Configure neighbor 10.10.10.4 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE1(config-router)#neighbor 10.10.10.2 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE1(config-router)#neighbor 10.10.10.3 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE1(config-router)#neighbor 10.10.10.4 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE1(config-router)#neighbor 10.10.10.2 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE1(config-router)#neighbor 10.10.10.3 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE1(config-router)#neighbor 10.10.10.4 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE1(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE1(config-router-af)#neighbor 10.10.10.2 activate | Activate EVPN for iBGP neighbor 10.10.10.2 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE1(config-router-af)#neighbor 10.10.10.3 activate | Activate EVPN for iBGP neighbor 10.10.10.3 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE1(config-router-af)#neighbor 10.10.10.4 activate | Activate EVPN for iBGP neighbor 10.10.10.4 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE1(config-router-af)#exit | Exit address family mode and return to the router BGP mode. |

| | |
|---------------------------|--|
| PE1(config-router)#commit | Commit the transaction. |
| PE1(config-router)#exit | Exit router BGP mode and return to the configure mode. |

PE1: MAC VRF Configuration

The below MAC VRF configuration on PE1 is carried out to define and set up VRFs named `vrf2` and `vppls1001` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|---|---|
| PE1(config)#mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE1(config-vrf)#rd 10.10.10.1:1700 | Configure Route-Distinguisher value of 10.10.10.1:1700. |
| PE1(config-vrf)#route-target both 1700:1700 | Configure import and export values for the <code>vrf2</code> as 1700:1700. |
| PE1(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE1(config)#mac vrf vppls1001 | Enter VRF mode named <code>vppls1001</code> . |
| PE1(config-vrf)#rd 10.10.10.1:1001 | Configure Route-Distinguisher value of 10.10.10.1:1001. |
| PE1(config-vrf)#route-target both 1001:1001 | Configure import and export values for the <code>vppls1001</code> as 1001:1001. |
| PE1(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: EVPN and VRF Mapping

The below EVPN and VRF mapping configuration on PE1 is performed to establish mappings between EVPN identifiers and VRFs, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|---|--|
| PE1(config)#evpn mpls id 1800 xconnect target-mpls-id 1700 | Configure the EVPN-VPWS identifier with a source identifier of 1800 and a target identifier of 1700. |
| PE1(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Map VRF <code>vrf2</code> to the EVPN-VPWS identifier |
| PE1(config-evpn-mpls)#commit | Commit the transaction. |
| PE1(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| PE1(config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier an identifier of 3000. |
| PE1(config-evpn-mpls)#host-reachability-protocol evpn-bgp vppls1001 | Map VRF <code>vppls1001</code> to the EVPN-VPWS identifier |
| PE1(config-evpn-mpls)#commit | Commit the transaction. |
| PE1(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE1: Access Port Configuration for Port-active

The below access port configuration for port-active mode on PE1 is carried out to configure various parameters including system-mac, load balancing, service carving preferences, and EVPN settings for efficient network access and connectivity.

| | |
|--|--|
| PE1(config)#interface po1 | Enter the port channel interface mode for po1 |
| PE1(config-if)#load-interval 30 | Set the load interval to 30. |
| PE1(config-if)#evpn multi-homed system-mac 0000.1111.7777 load-balancing port-active | Configure the system-mac address 0000.1111.7777 for port-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE1(config-if-es)#service-carving auto | Configure service carving as auto. |
| PE1(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE1(config-if)#exit | Exit interface mode po1 and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |
| PE1(config)#interface po1.1 switchport | Create a Layer 2 sub-interface po1.1 within the port channel. |
| PE1(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE1(config-if)#load-interval 30 | Set the load interval to 30. |
| PE1(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE1(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE1(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE1(config-if)#exit | Exit interface mode po1.1 and return to the configure mode. |
| PE1(config)#interface xe0 | Enter the interface mode for xe0. |
| PE1(config-if)#speed 10g | Set the speed to 10g. |
| PE1(config-if)#channel-group 1 mode active | Attach LAG interface po1. |
| PE1(config-if)#exit | Exit interface mode xe0 and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: Access Port Configuration for Single-active

The below access port configuration for single-active mode on PE1 is implemented to set up various parameters, including Ethernet Segment Identifier (ESI) settings, service carving preferences, and EVPN configurations, ensuring efficient network access and connectivity.

| | |
|---------------------------|--|
| PE1(config)#interface sa1 | Enter the single active interface mode for sa1 |
|---------------------------|--|

| | |
|---|--|
| PE1(config-if)#load-interval 30 | Set the load interval to 30. |
| PE1(config-if)#evpn multi-homed esi 00:00:11:11:77:77 load-balancing single-active | Configure the ESI with the value with the value 00:00:11:11:77:77 for single-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE1(config-if-es)#service-carving auto | Configure service carving as <code>auto</code> , allowing automatic determination of service distribution preferences. |
| PE1(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE1(config-if)#exit | Exit interface mode <code>sa1</code> and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |
| PE1(config)#interface sa1.1 switchport | Create a Layer 2 sub-interface <code>sa1.1</code> within the port channel. |
| PE1(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE1(config-if)#load-interval 30 | Set the load interval to 30. |
| PE1(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE1(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE1(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE1(config-if)#exit | Exit interface mode <code>sa1.1</code> and return to the configure mode. |
| PE1(config)#interface xe0 | Enter the interface mode for <code>xe0</code> . |
| PE1(config-if)#speed 10g | Set the speed to 10g. |
| PE1(config-if)#static-channel-group 1 | Attach the static-channel-group 1, the LAG interface <code>sa1</code> to <code>xe0</code> . |
| PE1(config-if)#exit | Exit interface mode <code>xe0</code> and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE2 Router Configurations

PE2: Loopback Interface

The configuration on PE2 for a loopback interface with IP address `10.10.10.2/32` secondary is set up to provide IP connectivity for the router.

| | |
|---|---|
| PE2#configure terminal | Enter configure mode. |
| PE2(config)#interface lo | Enter the interface mode for the loopback interface <code>lo</code> . |
| PE2(config-if)#ip address 10.10.10.2/32 secondary | Configure a secondary IP address, |

| | |
|---------------------|---|
| | 10.10.10.2/32, on the loopback interface. |
| PE2(config-if)#exit | Exit interface mode lo. |
| PE2(config)#commit | Commit the transaction. |

PE2: Global LDP

The configuration on PE2 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|--|---|
| PE2(config)#router ldp | Enter the Router LDP mode. |
| PE2(config-router)#router-id 10.10.10.2 | Set the router ID for LDP to 10.10.10.2. |
| PE2(config-router)#transport-address ipv4 10.10.10.2 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| PE2(config-router)#targeted-peer ipv4 10.10.10.1 | Configure targeted peer for LDP using IPv4 addresses. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE2(config-router)#targeted-peer ipv4 10.10.10.3 | Configure targeted peer for LDP using IPv4 addresses. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE2(config-router)#targeted-peer ipv4 10.10.10.4 | Configure targeted peer for LDP using IPv4 addresses. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE2(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: Global EVPN MPLS Command

The configuration on PE2 for the Global EVPN MPLS, includes activating EVPN MPLS, defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multi-homing, and activating EVPN MPLS multi-homing functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|--|--|
| PE2(config)#evpn mpls enable | Activate the EVPN MPLS functionality on PE2, enabling it to participate in EVPN MPLS services. |
| PE2(config)#commit | Commit candidate configuration to be running configuration. |
| PE2(config)#evpn mpls vtep-ip-global 10.10.10.2 | Configure the global VTEP IP address 10.10.10.2, associating it with the loopback IP. |
| PE2(config)#hardware-profile filter evpn-mpls-mh | Enable hardware-profile filter for EVPN MPLS multi- |

| | |
|--|--|
| enable | homing. |
| PE2(config)#evpn mpls multihoming enable | Activate the EVPN MPLS multi-homing functionality, allowing PE2 to support multi-homed EVPN MPLS services. |
| PE2(config)#commit | Commit the transaction. |

PE2: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE2 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|---------------------------------------|--|
| PE2(config)#interface xe4 | Enter interface mode xe4. |
| PE2(config-if)#ip address 10.1.3.1/30 | Configure an IP address, 10.1.3.1/30, on the interface xe4. |
| PE2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE2(config-if)#exit | Exit interface mode xe4. |
| PE2(config)#commit | Commit the transaction. |
| PE2(config)#interface xe5 | Enter interface mode xe5. |
| PE2(config-if)#ip address 10.1.4.1/30 | Configure an IP address, 10.1.4.1/30, on the interface xe5. |
| PE2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE2(config-if)#exit | Exit interface mode xe5. |
| PE2(config)#commit | Commit the transaction. |

PE2: OSPF Configuration

The below configuration is performed to set up OSPF on PE2, specifying the router ID, defining network interfaces, and configuring BFD parameters for efficient routing.

| | |
|---|--|
| PE2(config)#router ospf 100 | Enter the router OSPF mode. Configure PE2 to run OSPF with process ID 100. |
| PE2(config-router)#ospf router-id 10.10.10.2 | Set the OSPF router ID to 10.10.10.2, identifying PE2 within the OSPF network. |
| PE2(config-router)#network 10.1.3.1/30 area 0.0.0.0 | Advertise loopback address in OSPF. |
| PE2(config-router)#network 10.1.4.1/30 area 0.0.0.0 | Advertise network address in OSPF. |
| PE2(config-router)#bfd interval 3 minrx 3 | Configure BFD interval with an interval of 3, a |

| | |
|--------------------------------------|---|
| <code>multiplier 3</code> | minimum receive interval of 3, and a multiplier of 3. |
| <code>PE2(config-router)#exit</code> | Exit router OSPF mode and return to the configure mode. |
| <code>PE2(config)#commit</code> | Commit the transaction. |

PE2: BGP Configuration

The below BGP configuration on PE2 is established to enable BGP routing with ASN 65010, set the BGP router ID, define iBGP neighbors, configure BFD, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|--|--|
| <code>PE2(config)#router bgp 65010</code> | Enter the Router BGP mode, ASN: 65010 |
| <code>PE2(config-router)#bgp router-id 10.10.10.2</code> | Configure BGP router ID 10.10.10.2, identifying PE2 within the BGP network. |
| <code>PE2(config-router)#neighbor 10.10.10.1 remote-as 65010</code> | Configure neighbor 10.10.10.1 as an iBGP neighbor with their remote AS number 65010. |
| <code>PE2(config-router)#neighbor 10.10.10.1 update-source lo</code> | Configure neighbor 10.10.10.1 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| <code>PE2(config-router)#neighbor 10.10.10.3 remote-as 65010</code> | Configure neighbor 10.10.10.3 as an iBGP neighbor with their remote AS number 65010. |
| <code>PE2(config-router)#neighbor 10.10.10.3 update-source lo</code> | Configure neighbor 10.10.10.3 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| <code>PE2(config-router)#neighbor 10.10.10.4 remote-as 65010</code> | Configure neighbor 10.10.10.4 as an iBGP neighbor with their remote AS number 65010. |
| <code>PE2(config-router)#neighbor 10.10.10.4 update-source lo</code> | Configure neighbor 10.10.10.4 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| <code>PE2(config-router)#neighbor 10.10.10.1 fall-over bfd multihop</code> | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| <code>PE2(config-router)#neighbor 10.10.10.3 fall-over bfd multihop</code> | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| <code>PE2(config-router)#neighbor 10.10.10.4 fall-over bfd multihop</code> | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| <code>PE2(config-router)#neighbor 10.10.10.1 advertisement-interval 0</code> | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| <code>PE2(config-router)#neighbor 10.10.10.3 advertisement-interval 0</code> | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| <code>PE2(config-router)#neighbor 10.10.10.4 advertisement-interval 0</code> | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| <code>PE2(config-router)#address-family l2vpn evpn</code> | Enter into address family mode for L2VPN EVPN. |
| <code>PE2(config-router-af)#neighbor 10.10.10.1 activate</code> | Activate EVPN for iBGP neighbor 10.10.10.1 |

| | |
|--|---|
| | within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE2(config-router-af)#neighbor 10.10.10.3 activate | Activate EVPN for iBGP neighbor 10.10.10.3 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE2(config-router-af)#neighbor 10.10.10.4 activate | Activate EVPN for iBGP neighbor 10.10.10.4 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE2(config-router-af)#exit | Exit address family mode and return to the router BGP mode. |
| PE2(config-router)#commit | Commit the transaction. |
| PE2(config-router)#exit | Exit router BGP mode and return to the configure mode. |

PE2: MAC VRF Configuration

The below MAC VRF configuration on PE2 is carried out to define and set up VRFs named `vrf2` and `vp1s1001` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|---|--|
| PE2(config)#mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE2(config-vrf)#rd 10.10.10.2:1700 | Configure Route-Distinguisher value of 10.10.10.2:1700. |
| PE2(config-vrf)#route-target both 1700:1700 | Configure import and export values for the <code>vrf2</code> as 1700:1700. |
| PE2(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE2(config)#mac vrf vp1s1001 | Enter VRF mode named <code>vp1s1001</code> . |
| PE2(config-vrf)#rd 10.10.10.2:1001 | Configure Route-Distinguisher value of 10.10.10.2:1001. |
| PE2(config-vrf)#route-target both 1001:1001 | Configure import and export values for the <code>vp1s1001</code> as 1001:1001. |
| PE2(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: EVPN and VRF Mapping

The below EVPN and VRF mapping configuration on PE2 is performed to establish mappings between EVPN identifiers and VRFs, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|--|--|
| PE2(config)#evpn mpls id 1800 xconnect target-mpls-id 1700 | Configure the EVPN-VPWS identifier with a source identifier of 1800 and a target identifier of 1700. |
| PE2(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Map VRF <code>vrf2</code> to the EVPN-VPWS identifier |
| PE2(config-evpn-mpls)#commit | Commit the transaction. |

| | |
|--|---|
| PE2(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| PE2(config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier an identifier of 3000. |
| PE2(config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Map VRF vpls1001 to the EVPN-VPWS identifier |
| PE2(config-evpn-mpls)#commit | Commit the transaction. |
| PE2(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: Access Port Configuration for Port-active

The below access port configuration for port-active mode on PE2 is carried out to configure various parameters including system-mac, load balancing, service carving preferences, and EVPN settings for efficient network access and connectivity.

| | |
|--|--|
| PE2(config)#interface po1 | Enter the port channel interface mode for po1 |
| PE2(config-if)#load-interval 30 | Set the load interval to 30. |
| PE2(config-if)#evpn multi-homed system-mac 0000.1111.7777 load-balancing port-active | Configure the system-mac address 0000.1111.7777 for port-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE2(config-if-es)#service-carving auto | Configure service carving as auto, allowing automatic determination of service distribution preferences. |
| PE2(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE2(config-if)#exit | Exit interface mode po1 and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |
| PE2(config)#interface po1.1 switchport | Create a Layer 2 sub-interface po1.1 within the port channel. |
| PE2(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE2(config-if)#load-interval 30 | Set the load interval to 30. |
| PE2(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE2(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE2(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE2(config-if)#exit | Exit interface mode po1.1 and return to the configure mode. |
| PE2(config)#interface xe08 | Enter the interface mode for xe8. |
| PE2(config-if)#speed 10g | Set the speed to 10g. |

| | |
|--|---|
| PE2(config-if)#channel-group 1 mode active | Attach LAG interface po1. |
| PE2(config-if)#exit | Exit interface mode xe8 and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: Access Port Configuration for Single-active

The below access port configuration for single-active mode on PE2 is implemented to set up various parameters, including Ethernet Segment Identifier (ESI) settings, service carving preferences, and EVPN configurations, ensuring efficient network access and connectivity.

| | |
|---|---|
| PE2(config)#interface sa2 | Enter the single active interface mode for sa2. |
| PE2(config-if)#load-interval 30 | Set the load interval to 30. |
| PE2(config-if)#evpn multi-homed esi 00:00:11:11:77:77 load-balancing single-active | Configure the ESI with the value 00:00:11:11:77:77 for single-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE2(config-if-es)#service-carving auto | Configure service carving as auto, allowing automatic determination of service distribution preferences. |
| PE2(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE2(config-if)#exit | Exit interface mode sa2 and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |
| PE2(config)#interface sa2.1 switchport | Create a Layer 2 sub-interface sa2.1 within the port channel. |
| PE2(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE2(config-if)#load-interval 30 | Set the load interval to 30. |
| PE2(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE2(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE2(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE2(config-if)#exit | Exit interface mode sa2.1 and return to the configure mode. |
| PE2(config)#interface xe8 | Enter the interface mode for xe8. |
| PE2(config-if)#speed 10g | Set the speed to 10g. |
| PE2(config-if)#static-channel-group 2 | Attach the static-channel-group 2, the LAG interface sa2 to xe8. |
| PE2(config-if)#exit | Exit interface mode xe8 and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE3 Router Configurations

PE3: Loopback Interface

The configuration on PE3 for a loopback interface with IP address 10.10.10.3/32 secondary is set up to provide IP connectivity for the router.

| | |
|---|---|
| PE3#configure terminal | Enter configure mode. |
| PE3(config)#interface lo | Enter the interface mode for the loopback interface lo. |
| PE3(config-if)#ip address 10.10.10.3/32 secondary | Configure a secondary IP address, 10.10.10.3/32, on the loopback interface. |
| PE3(config-if)#exit | Exit interface mode lo. |
| PE3(config)#commit | Commit the transaction. |

PE3: Global LDP

The configuration on PE3 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|--|---|
| PE3(config)#router ldp | Enter the Router LDP mode. |
| PE3(config-router)#router-id 10.10.10.3 | Set the router ID for LDP to 10.10.10.3. |
| PE3(config-router)#transport-address ipv4 10.10.10.3 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| PE3(config-router)#targeted-peer ipv4 10.10.10.1 | Configure targeted peer for LDP using IPv4 addresses. |
| PE3(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE3(config-router)#targeted-peer ipv4 10.10.10.2 | Configure targeted peer for LDP using IPv4 addresses. |
| PE3(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE3(config-router)#targeted-peer ipv4 10.10.10.4 | Configure targeted peer for LDP using IPv4 addresses. |
| PE3(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE3(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: Global EVPN MPLS Command

The configuration on PE3 for the Global EVPN MPLS, includes activating EVPN MPLS, defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multi-homing, and activating EVPN MPLS multi-homing

functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|---|--|
| PE3(config)#evpn mpls enable | Activate the EVPN MPLS functionality on PE3, enabling it to participate in EVPN MPLS services. |
| PE3(config)#commit | Commit candidate configuration to be running configuration. |
| PE3(config)#evpn mpls vtep-ip-global 10.10.10.3 | Configure the global VTEP IP address 10.10.10.3, associating it with the loopback IP. |
| PE3(config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter for EVPN MPLS multi-homing. |
| PE3(config)#evpn mpls multihoming enable | Activate the EVPN MPLS multi-homing functionality, allowing PE3 to support multi-homed EVPN MPLS services. |
| PE3(config)#commit | Commit the transaction. |

PE3: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE3 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|---------------------------------------|--|
| PE3(config)#interface xe1 | Enter interface mode xe1. |
| PE3(config-if)#ip address 10.1.6.1/30 | Configure an IP address, 10.1.6.1/30, on the interface xe1. |
| PE3(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE3(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE3(config-if)#exit | Exit interface mode xe1. |
| PE3(config)#commit | Commit the transaction. |
| PE3(config)#interface xe5 | Enter interface mode xe5. |
| PE3(config-if)#ip address 10.1.7.1/30 | Configure an IP address, 10.1.7.1/30, on the interface xe5. |
| PE3(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE3(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE3(config-if)#exit | Exit interface mode xe5. |
| PE3(config)#commit | Commit the transaction. |

PE3: OSPF Configuration

The below configuration is performed to set up OSPF on PE3, specifying the router ID, defining network interfaces, and configuring BFD parameters for efficient routing.

| | |
|--|---|
| PE3(config)#router ospf 100 | Enter the router OSPF mode. Configure PE3 to run OSPF with process ID 100. |
| PE3(config-router)#ospf router-id 10.10.10.3 | Set the OSPF router ID to 10.10.10.3, identifying PE3 within the OSPF network. |
| PE3(config-router)#network 10.10.10.3/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| PE3(config-router)#network 10.1.6.1/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| PE3(config-router)#network 10.1.7.1/30 area 0.0.0.0 | Advertise network address in OSPF. |
| PE3(config-router)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval with an interval of 3, a minimum receive interval of 3, and a multiplier of 3. |
| PE3(config-router)#exit | Exit router OSPF mode and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: BGP Configuration

The below BGP configuration on PE3 is established to enable BGP routing with ASN 65010, set the BGP router ID, define iBGP neighbors, configure BFD, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|---|--|
| PE3(config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| PE3(config-router)#bgp router-id 10.10.10.3 | Configure BGP router ID 10.10.10.2, identifying PE3 within the BGP network. |
| PE3(config-router)#neighbor 10.10.10.1 remote-as 65010 | Configure neighbor 10.10.10.1 as an iBGP neighbor with their remote AS number 65010. |
| PE3(config-router)#neighbor 10.10.10.1 update-source lo | Configure neighbor 10.10.10.1 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE3(config-router)#neighbor 10.10.10.2 remote-as 65010 | Configure neighbor 10.10.10.2 as an iBGP neighbor with their remote AS number 65010. |
| PE3(config-router)#neighbor 10.10.10.2 update-source lo | Configure neighbor 10.10.10.2 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE3(config-router)#neighbor 10.10.10.4 remote-as 65010 | Configure neighbor 10.10.10.4 as an iBGP neighbor with their remote AS number 65010. |
| PE3(config-router)#neighbor 10.10.10.4 update-source lo | Configure neighbor 10.10.10.4 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE3(config-router)#neighbor 10.10.10.1 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE3(config-router)#neighbor 10.10.10.2 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE3(config-router)#neighbor 10.10.10.4 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |

| | |
|---|---|
| PE3(config-router)#neighbor 10.10.10.1 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE3(config-router)#neighbor 10.10.10.2 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE3(config-router)#neighbor 10.10.10.4 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE3(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE3(config-router-af)#neighbor 10.10.10.1 activate | Activate EVPN for iBGP neighbor 10.10.10.1 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE3(config-router-af)#neighbor 10.10.10.2 activate | Activate EVPN for iBGP neighbor 10.10.10.2 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE3(config-router-af)#neighbor 10.10.10.4 activate | Activate EVPN for iBGP neighbor 10.10.10.4 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE3(config-router-af)#exit | Exit address family mode and return to the router BGP mode. |
| PE3(config-router)#commit | Commit the transaction. |
| PE3(config-router)#exit | Exit router BGP mode and return to the configure mode. |

PE3: MAC VRF Configuration

The below MAC VRF configuration on PE3 is carried out to define and set up VRFs named `vrf2` and `vp1s1001` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|---|--|
| PE3(config)#mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE3(config-vrf)#rd 10.10.10.3:1700 | Configure Route-Distinguisher value of 10.10.10.3:1700. |
| PE3(config-vrf)#route-target both 1700:1700 | Configure import and export values for the <code>vrf2</code> as 1700:1700. |
| PE3(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE3(config)#mac vrf vp1s1001 | Enter VRF mode named <code>vp1s1001</code> . |
| PE3(config-vrf)#rd 10.10.10.3:1001 | Configure Route-Distinguisher value of 10.10.10.3:1001. |
| PE3(config-vrf)#route-target both 1001:1001 | Configure import and export values for the <code>vp1s1001</code> as 1001:1001. |
| PE3(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: EVPN and VRF Mapping

The below EVPN and VRF mapping configuration on PE3 is performed to establish mappings between EVPN identifiers and VRFs, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|--|--|
| PE3(config)#evpn mpls id 1700 xconnect target-mpls-id 1800 | Configure the EVPN-VPWS identifier with a source identifier of 1700 and a target identifier of 1800. |
| PE3(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Map VRF <code>vrf2</code> to the EVPN-VPWS identifier |
| PE3(config-evpn-mpls)#commit | Commit the transaction. |
| PE3(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| PE3(config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier an identifier of 3000. |
| PE3(config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Map VRF <code>vpls1001</code> to the EVPN-VPWS identifier |
| PE3(config-evpn-mpls)#commit | Commit the transaction. |
| PE3(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE3: Access Port Configuration for Port-active

The below access port configuration for port-active mode on PE3 is carried out to configure various parameters including system-MAC, load balancing, service carving preferences, and EVPN settings for efficient network access and connectivity.

| | |
|--|--|
| PE3(config)#interface po1 | Enter the port channel interface mode for <code>po1</code> |
| PE3(config-if)#load-interval 30 | Set the load interval to 30. |
| PE3(config-if)#evpn multi-homed system-mac 0000.2222.7777 load-balancing port-active | Configure the system-mac address 0000.2222.7777 for port-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE3(config-if-es)#service-carving auto | Configure service carving as <code>auto</code> , allowing automatic determination of service distribution preferences. |
| PE3(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE3(config-if)#exit | Exit interface mode <code>po1</code> and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |
| PE3(config)#interface po1.1 switchport | Create a Layer 2 sub-interface <code>po1.1</code> within the port channel. |
| PE3(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE3(config-if)#load-interval 30 | Set the load interval to 30. |
| PE3(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID |

| | |
|--|--|
| | configuration. |
| PE3(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE3(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE3(config-if)#exit | Exit interface mode <code>po1.1</code> and return to the configure mode. |
| PE3(config)#interface xe2 | Enter the interface mode for <code>xe2</code> . |
| PE3(config-if)#speed 10g | Set the speed to 10g. |
| PE3(config-if)#channel-group 1 mode active | Attach LAG interface <code>po1</code> . |
| PE3(config-if)#exit | Exit interface mode <code>xe2</code> and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: Access Port Configuration for Single-active

The below access port configuration for single-active mode on PE3 is implemented to set up various parameters, including Ethernet Segment Identifier (ESI) settings, service carving preferences, and EVPN configurations, ensuring efficient network access and connectivity.

| | |
|---|--|
| PE3(config)#interface sa1 | Enter the single active interface mode for <code>sa1</code> . |
| PE3(config-if)#load-interval 30 | Set the load interval to 30. |
| PE3(config-if)#evpn multi-homed esi 00:00:22:22:77:77 load-balancing single-active | Configure the ESI with the value with the value 00:00:22:22:77:77 for single-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE3(config-if-es)#service-carving auto | Configure service carving as <code>auto</code> , allowing automatic determination of service distribution preferences. |
| PE3(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE3(config-if)#exit | Exit interface mode <code>sa1</code> and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |
| PE3(config)#interface sa1.1 switchport | Create a Layer 2 sub-interface <code>sa1.1</code> within the port channel. |
| PE3(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE3(config-if)#load-interval 30 | Set the load interval to 30. |
| PE3(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE3(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE3(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |

| | |
|---------------------------------------|---|
| PE3(config-if)#exit | Exit interface mode <code>sa1.1</code> and return to the configure mode. |
| PE3(config)#interface xe2 | Enter the interface mode for <code>xe2</code> . |
| PE3(config-if)#speed 10g | Set the speed to 10g. |
| PE3(config-if)#static-channel-group 1 | Attach the static-channel-group 1, the LAG interface <code>sa1</code> to <code>xe2</code> . |
| PE3(config-if)#exit | Exit interface mode <code>xe2</code> and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE4 Router Configurations

PE4: Loopback Interface

The configuration on PE4 for a loopback interface with IP address `10.10.10.4/32` secondary is set up to provide IP connectivity for the router.

| | |
|---|---|
| PE4#configure terminal | Enter configure mode. |
| PE4(config)#interface lo | Enter the interface mode for the loopback interface <code>lo</code> . |
| PE4(config-if)#ip address 10.10.10.4/32 secondary | Configure a secondary IP address, <code>10.10.10.4/32</code> , on the loopback interface. |
| PE4(config-if)#exit | Exit interface mode <code>lo</code> . |
| PE4(config)#commit | Commit the transaction. |

PE4: Global LDP

The configuration on PE4 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|--|--|
| PE4(config)#router ldp | Enter the Router LDP mode. |
| PE4(config-router)#router-id 10.10.10.4 | Set the router ID for LDP to <code>10.10.10.4</code> . |
| PE4(config-router)#transport-address ipv4 10.10.10.4 | Configure the transport address for IPv4 (for IPv6 use <code>ipv6</code> parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| PE4(config-router)#targeted-peer ipv4 10.10.10.1 | Configure targeted peer for LDP using IPv4 addresses. |
| PE4(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE4(config-router)#targeted-peer ipv4 10.10.10.2 | Configure targeted peer for LDP using IPv4 addresses. |
| PE4(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |

| | |
|--|--|
| PE4(config-router)#targeted-peer ipv4 10.10.10.3 | Configure targeted peer for LDP using IPv4 addresses. |
| PE4(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE4(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |

PE4: Global EVPN MPLS Command

The configuration on PE4 for the Global EVPN MPLS, includes activating EVPN MPLS, defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multi-homing, and activating EVPN MPLS multi-homing functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|---|--|
| PE4(config)#evpn mpls enable | Activate the EVPN MPLS functionality on PE4, enabling it to participate in EVPN MPLS services. |
| PE4(config)#commit | Commit candidate configuration to be running configuration. |
| PE4(config)#evpn mpls vtep-ip-global 10.10.10.4 | Configure the global VTEP IP address 10.10.10.4, associating it with the loopback IP. |
| PE4(config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter for EVPN MPLS multi-homing. |
| PE4(config)#evpn mpls multihoming enable | Activate the EVPN MPLS multi-homing functionality, allowing PE4 to support multi-homed EVPN MPLS services. |
| PE4(config)#commit | Commit the transaction. |

PE4: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE4 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|---------------------------------------|--|
| PE4(config)#interface xe2 | Enter interface mode xe2. |
| PE4(config-if)#ip address 10.1.5.1/30 | Configure an IP address, 10.1.5.1/30, on the interface xe2. |
| PE4(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE4(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE4(config-if)#exit | Exit interface mode xe2. |
| PE4(config)#commit | Commit the transaction. |
| PE4(config)#interface xe0 | Enter interface mode xe0. |
| PE4(config-if)#ip address 10.1.8.1/30 | Configure an IP address, 10.1.8.1/30, on the interface xe0. |

| | |
|--------------------------------|--|
| PE4(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE4(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE4(config-if)#exit | Exit interface mode <code>xe0</code> . |
| PE4(config)#commit | Commit the transaction. |

PE4: OSPF Configuration

The below configuration is performed to set up OSPF on PE4, specifying the router ID, defining network interfaces, and configuring BFD parameters for efficient routing.

| | |
|--|---|
| PE4(config)#router ospf 100 | Enter the router OSPF mode. Configure PE4 to run OSPF with process ID 100. |
| PE4(config-router)#ospf router-id 10.10.10.4 | Set the OSPF router ID to <code>10.10.10.4</code> , identifying PE3 within the OSPF network. |
| PE4(config-router)#network 10.10.10.4/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| PE4(config-router)#network 10.1.5.1/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| PE4(config-router)#network 10.1.8.1/30 area 0.0.0.0 | Advertise network address in OSPF. |
| PE4(config-router)#bfd interval 3 minrx 3 multiplier 3 | Configure BFD interval with an interval of 3, a minimum receive interval of 3, and a multiplier of 3. |
| PE4(config-router)#exit | Exit router OSPF mode and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |

PE4: BGP Configuration

The below BGP configuration on PE4 is established to enable BGP routing with ASN 65010, set the BGP router ID, define iBGP neighbors, configure BFD, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|---|--|
| PE4(config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| PE4(config-router)#bgp router-id 10.10.10.4 | Configure BGP router ID <code>10.10.10.4</code> , identifying PE4 within the BGP network. |
| PE4(config-router)#neighbor 10.10.10.1 remote-as 65010 | Configure neighbor 10.10.10.1 as an iBGP neighbor with their remote AS number 65010. |
| PE4(config-router)#neighbor 10.10.10.1 update-source lo | Configure neighbor 10.10.10.1 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE4(config-router)#neighbor 10.10.10.2 remote-as 65010 | Configure neighbor 10.10.10.2 as an iBGP neighbor with their remote AS number 65010. |
| PE4(config-router)#neighbor 10.10.10.2 update-source lo | Configure neighbor 10.10.10.2 as an iBGP neighbor, specifying the source of routing updates as the |

| | |
|---|---|
| | loopback interface. |
| PE4(config-router)#neighbor 10.10.10.3 remote-as 65010 | Configure neighbor 10.10.10.3 as an iBGP neighbor with their remote AS number 65010. |
| PE4(config-router)#neighbor 10.10.10.3 update-source lo | Configure neighbor 10.10.10.3 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE4(config-router)#neighbor 10.10.10.1 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE4(config-router)#neighbor 10.10.10.2 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE4(config-router)#neighbor 10.10.10.3 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE4(config-router)#neighbor 10.10.10.1 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE4(config-router)#neighbor 10.10.10.2 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE4(config-router)#neighbor 10.10.10.3 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE4(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE4(config-router-af)#neighbor 10.10.10.1 activate | Activate EVPN for iBGP neighbor 10.10.10.1 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE4(config-router-af)#neighbor 10.10.10.2 activate | Activate EVPN for iBGP neighbor 10.10.10.2 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE4(config-router-af)#neighbor 10.10.10.3 activate | Activate EVPN for iBGP neighbor 10.10.10.3 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE4(config-router-af)#exit | Exit address family mode and return to the router BGP mode. |
| PE4(config-router)#commit | Commit the transaction. |
| PE4(config-router)#exit | Exit router BGP mode and return to the configure mode. |

PE4: MAC VRF Configuration

The below MAC VRF configuration on PE4 is carried out to define and set up VRFs named `vrf2` and `vp1s1001` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|------------------------------------|---|
| PE4(config)#mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE4(config-vrf)#rd 10.10.10.4:1700 | Configure Route-Distinguisher value of 10.10.10.4:1700. |

| | |
|---|---|
| PE4(config-vrf)#route-target both 1700:1700 | Configure import and export values for the vrf2 as 1700:1700. |
| PE4(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE4(config)#mac vrf vpls1001 | Enter VRF mode named vpls1001. |
| PE4(config-vrf)#rd 10.10.10.4:1001 | Configure Route-Distinguisher value of 10.10.10.4:1001. |
| PE4(config-vrf)#route-target both 1001:1001 | Configure import and export values for the vpls1001 as 1001:1001. |
| PE4(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |

PE4: EVPN and VRF Mapping

The below EVPN and VRF mapping configuration on PE4 is performed to establish mappings between EVPN identifiers and VRFs, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|--|--|
| PE4(config)#evpn mpls id 1700 xconnect target-mpls-id 1800 | Configure the EVPN-VPWS identifier with a source identifier of 1700 and a target identifier of 1800. |
| PE4(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Map VRF vrf2 to the EVPN-VPWS identifier |
| PE4(config-evpn-mpls)#commit | Commit the transaction. |
| PE4(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| PE4(config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier an identifier of 3000. |
| PE4(config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Map VRF vpls1001 to the EVPN-VPWS identifier |
| PE4(config-evpn-mpls)#commit | Commit the transaction. |
| PE4(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE4: Access Port Configuration for Port-active

The below access port configuration for port-active mode on PE4 is carried out to configure various parameters including system-MAC, load balancing, service carving preferences, and EVPN settings for efficient network access and connectivity.

| | |
|--|--|
| PE4(config)#interface po1 | Enter the port channel interface mode for po1 |
| PE4(config-if)#load-interval 30 | Set the load interval to 30. |
| PE4(config-if)#evpn multi-homed system-mac 0000.2222.7777 load-balancing port-active | Configure the system-mac address 0000.2222.7777 for port-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE4(config-if-es)#service-carving auto | Configure service carving as auto, allowing automatic determination of service distribution |

| | |
|--|--|
| | preferences. |
| PE4(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE4(config-if)#exit | Exit interface mode <code>po1</code> and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |
| PE4(config)#interface po1.1 switchport | Create a Layer 2 sub-interface <code>po1.1</code> within the port channel. |
| PE4(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE4(config-if)#load-interval 30 | Set the load interval to 30. |
| PE4(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE4(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE4(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE4(config-if)#exit | Exit interface mode <code>po1.1</code> and return to the configure mode. |
| PE4(config)#interface xe11 | Enter the interface mode for <code>xe11</code> . |
| PE4(config-if)#speed 10g | Set the speed to 10g. |
| PE4(config-if)#channel-group 1 mode active | Attach LAG interface <code>po1</code> . |
| PE4(config-if)#exit | Exit interface mode <code>xe11</code> and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |

PE4: Access Port Configuration for Single-active

The below access port configuration for single-active mode on PE4 is implemented to set up various parameters, including Ethernet Segment Identifier (ESI) settings, service carving preferences, and EVPN configurations, ensuring efficient network access and connectivity.

| | |
|---|---|
| PE4(config)#interface sa2 | Enter the single active interface mode for <code>sa2</code> . |
| PE4(config-if)#load-interval 30 | Set the load interval to 30. |
| PE4(config-if)#evpn multi-homed esi 00:00:22:22:77:77 load-balancing single-active | Configure the ESI with the value with the value <code>00:00:22:22:77:77</code> for single-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE4(config-if-es)#service-carving auto | Configure service carving as <code>auto</code> , allowing automatic determination of service distribution preferences. |
| PE4(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE4(config-if)#exit | Exit interface mode <code>sa2</code> and return to the configure |

| | |
|---|---|
| | mode. |
| PE4(config)#commit | Commit the transaction. |
| PE4(config)#interface sa2.1 switchport | Create a Layer 2 sub-interface sa2.1 within the port channel. |
| PE4(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE4(config-if)#load-interval 30 | Set the load interval to 30. |
| PE4(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE4(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE4(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE4(config-if)#exit | Exit interface mode sa2.1 and return to the configure mode. |
| PE4(config)#interface xe11 | Enter the interface mode for xe11. |
| PE4(config-if)#speed 10g | Set the speed to 10g. |
| PE4(config-if)#static-channel-group 2 | Attach the static-channel-group 2, the LAG interface sa2 to xe11. |
| PE4(config-if)#exit | Exit interface mode xe11 and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |

Core Routers Configurations

P1: Loopback Interface

The configuration on P1 for a loopback interface with IP address 10.10.10.5/32 secondary is set up to provide IP connectivity for the router.

| | |
|--|---|
| P1#configure terminal | Enter configure mode. |
| P1(config)#interface lo | Enter the interface mode for the loopback interface lo. |
| P1(config-if)#ip address 10.10.10.5/32 secondary | Configure a secondary IP address, 10.10.10.5/32, on the loopback interface. |
| P1(config-if)#exit | Exit interface mode lo. |
| P1(config)#commit | Commit the transaction. |

P1: Global LDP

The configuration on P1 for the Global LDP router, specifying router ID and targeted peer, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|--|--|
| P1(config)#router ldp | Enter the Router LDP mode. |
| P1(config-router)#router-id 10.10.10.5 | Set the router ID for LDP to 10.10.10.5. |

| | |
|--|---|
| P1(config-router)#transport-address ipv4 10.10.10.5 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| P1(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| P1(config)#commit | Commit the transaction. |

P1: Interface Configuration

The below configuration is performed to set up interfaces on P1 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|--------------------------------------|--|
| P1(config)#interface xe1 | Enter interface mode xe1. |
| P1(config-if)#ip address 10.1.1.2/30 | Configure an IP address, 10.1.1.2/30, on the interface xe1. |
| P1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P1(config-if)#exit | Exit interface mode xe1. |
| P1(config)#commit | Commit the transaction. |
| P1(config)#interface xe2 | Enter interface mode xe2. |
| P1(config-if)#ip address 10.1.4.2/30 | Configure an IP address, 10.1.4.2/30, on the interface xe2. |
| P1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P1(config-if)#exit | Exit interface mode xe2. |
| P1(config)#commit | Commit the transaction. |
| P1(config)#interface xe3 | Enter interface mode xe3. |
| P1(config-if)#ip address 10.1.5.2/30 | Configure an IP address, 10.1.5.2/30, on the interface xe3. |
| P1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P1(config-if)#exit | Exit interface mode xe3. |

| | |
|--------------------------------------|--|
| P1(config)#commit | Commit the transaction. |
| P1(config)#interface xe4 | Enter interface mode xe4. |
| P1(config-if)#ip address 10.1.6.2/30 | Configure an IP address, 10.1.6.2/30, on the interface xe4. |
| P1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P1(config-if)#exit | Exit interface mode xe4. |
| P1(config)#commit | Commit the transaction. |

P1: OSPF Configuration

The below configuration is performed to set up OSPF on P1, specifying the router ID, and defining network interfaces for efficient routing.

| | |
|--|---|
| P1(config)#router ospf 100 | Enter the router OSPF mode. Configure P1 to run OSPF with process ID 100. |
| P1(config-router)#ospf router-id 10.10.10.5 | Set the OSPF router ID to 10.10.10.5, identifying P1 within the OSPF network. |
| P1(config-router)#network 10.10.10.5/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| P1(config-router)#network 10.1.1.2/30 area 0.0.0.0 | Advertise network address in OSPF. |
| P1(config-router)#network 10.1.4.2/30 area 0.0.0.0 | Advertise network address in OSPF. |
| P1(config-router)#network 10.1.5.2/30 area 0.0.0.0 | Advertise network address in OSPF. |
| P1(config-router)#network 10.1.6.2/30 area 0.0.0.0 | Advertise network address in OSPF. |
| P1(config-router)#exit | Exit router OSPF mode and return to the configure mode. |
| P1(config)#commit | Commit the transaction. |

P2: Loopback Interface

The configuration on P2 for a loopback interface with IP address 10.10.10.6/32 secondary is set up to provide IP connectivity for the router.

| | |
|--|---|
| P2#configure terminal | Enter configure mode. |
| P2(config)#interface lo | Enter the interface mode for the loopback interface lo. |
| P2(config-if)#ip address 10.10.10.6/32 secondary | Configure a secondary IP address, 10.10.10.6/32, on the loopback interface. |
| P2(config-if)#exit | Exit interface mode lo. |
| P2(config)#commit | Commit the transaction. |

P2: Global LDP

The configuration on P2 for the Global LDP router, specifying router ID and targeted peer, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|---|---|
| P2(config)#router ldp | Enter the Router LDP mode. |
| P2(config-router)#router-id 10.10.10.6 | Set the router ID for LDP to 10.10.10.6. |
| P2(config-router)#transport-address ipv4 10.10.10.6 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| P2(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| P2(config)#commit | Commit the transaction. |

P2: Interface Configuration

The below configuration is performed to set up interfaces on P2 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|--------------------------------------|--|
| P2(config)#interface xe12 | Enter interface mode xe12. |
| P2(config-if)#ip address 10.1.2.2/30 | Configure an IP address, 10.1.2.2/30, on the interface xe12. |
| P2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P2(config-if)#exit | Exit interface mode xe12. |
| P2(config)#commit | Commit the transaction. |
| P2(config)#interface xe13 | Enter interface mode xe13. |
| P2(config-if)#ip address 10.1.3.2/30 | Configure an IP address, 10.1.3.2/30, on the interface xe13. |
| P2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P2(config-if)#exit | Exit interface mode xe13. |
| P2(config)#commit | Commit the transaction. |
| P2(config)#interface xe11 | Enter interface mode xe11. |
| P2(config-if)#ip address 10.1.7.2/30 | Configure an IP address, 10.1.7.2/30, on the interface xe11. |

| | |
|--------------------------------------|--|
| P2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P2(config-if)#exit | Exit interface mode <code>xe11</code> . |
| P2(config)#commit | Commit the transaction. |
| P2(config)#interface xe14 | Enter interface mode <code>xe14</code> . |
| P2(config-if)#ip address 10.1.8.2/30 | Configure an IP address, <code>10.1.8.2/30</code> , on the interface <code>xe14</code> . |
| P2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P2(config-if)#exit | Exit interface mode <code>xe14</code> . |
| P2(config)#commit | Commit the transaction. |

P2: OSPF Configuration

The below configuration is performed to set up OSPF on P2, specifying the router ID, and defining network interfaces for efficient routing.

| | |
|--|---|
| P2(config)#router ospf 100 | Enter the router OSPF mode. Configure P2 to run OSPF with process ID 100. |
| P2(config-router)#ospf router-id 10.10.10.6 | Set the OSPF router ID to <code>10.10.10.6</code> , identifying P2 within the OSPF network. |
| P2(config-router)#network 10.10.10.6/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| P2(config-router)#network 10.1.2.2/30 area 0.0.0.0 | Advertise network address in OSPF. |
| P2(config-router)#network 10.1.3.2/30 area 0.0.0.0 | Advertise network address in OSPF. |
| P2(config-router)#network 10.1.7.2/30 area 0.0.0.0 | Advertise network address in OSPF. |
| P2(config-router)#network 10.1.8.2/30 area 0.0.0.0 | Advertise network address in OSPF. |
| P2(config-router)#exit | Exit router OSPF mode and return to the configure mode. |
| P2(config)#commit | Commit the transaction. |

EVPN MPLS Active-Standby MH Validation

The following show outputs provide validation results for both single-active and port-active modes, covering ELINE and ELAN services configurations with LDP as the underlay MPLS path.

Single-Active

The following show output displays the types of load-balancing port selection criteria (PSC) used on configured static aggregators for CE1, PE1, PE2, PE3, PE4, and CE2 devices in the network [Figure 54](#) using the **show static-channel-group** command.

```
CE1#show static-channel-group
Static Aggregator: sa1
Member Status
xe48    up
-----

Static Aggregator: sa2
Member Status
xe50    up

PE1#show static-channel-group
Static Aggregator: sa1
Member Status weight
xe0     up

PE2#show static-channel-group
Static Aggregator: sa2
Member Status weight
xe8     up

PE3#show static-channel-group
Static Aggregator: sa1
Member Status weight
xe2     up
PE4#show static-channel-group
Static Aggregator: sa2
Member Status weight
xe11    up

CE2#show static-channel-group
Static Aggregator: sa1
Member Status weight
xe23    up
-----

Static Aggregator: sa2
Member Status weight
ge11    up
```

Single-Active ELINE

The following show output displays the active EVPN MPLS Tunnels and load balance for ELINE on PE1, PE2, PE3, and PE4 devices in the network [Figure 54](#) using the **show evpn load-balance all** and **show evpn mpls xconnect tunnel** commands.

```
PE1#show evpn load-balance all
ESI          AC-IF/PE    PE-IP-ADDRESS  Redundancy    Service-
carving      weight    Revertive      AC-DF    Status
=====
00:11:22:33:00:00:00:55:66:77 sa1.1        10.10.10.1     single-
active       auto      NO             0        NO        ACTIVE
00:11:22:33:00:00:00:55:66:77 ----         10.10.10.2     single-
active       auto      NO             0        NO        ----

PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination    Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====
===
```

```

10.10.10.1      10.10.10.4      Installed      00:14:05      00:03:58      1800          1700
10.10.10.1      10.10.10.3      Installed      00:14:05      00:04:29      1800          1700

```

Total number of entries are 2

PE2#show evpn load-balance all

| ESI | carving | weight | Revertive | AC-IF/PE AC-DF | PE-IP-ADDRESS Status | Redundancy | Service- |
|-------------------------------|---------|--------|-----------|-------------------|-------------------------|------------|----------|
| 00:11:22:33:00:00:00:55:66:77 | --- | | | | 10.10.10.1 | single- | |
| active | auto | | 0 | NO | NO | --- | |
| 00:11:22:33:00:00:00:55:66:77 | | | sa2.1 | | 10.10.10.2 | single- | |
| active | auto | | 0 | NO | NO | STANDBY | |

ELINE

PE2#show evpn mpls xconnect tunnel

| EVPN-MPLS Network tunnel Entries | | | | | | | |
|----------------------------------|-------------|-----------|----------|----------|---------------|---------|---------|
| Source | Destination | Status | Up/Down | Update | local-evpn-id | remote- | evpn-id |
| 10.10.10.2 | 10.10.10.4 | Installed | 00:12:33 | 00:04:08 | 1800 | 1700 | |
| 10.10.10.2 | 10.10.10.3 | Installed | 00:12:33 | 00:04:08 | 1800 | 1700 | |

Total number of entries are 2

PE3#show evpn load-balance all

| ESI | carving | weight | Revertive | AC-IF/PE AC-DF | PE-IP-ADDRESS Status | Redundancy | Service- |
|-------------------------------|---------|--------|-----------|-------------------|-------------------------|------------|----------|
| 00:12:22:33:00:00:00:55:66:77 | | | sa1.1 | | 10.10.10.3 | single- | |
| active | auto | | 0 | NO | NO | ACTIVE | |
| 00:12:22:33:00:00:00:55:66:77 | | | --- | | 10.10.10.4 | single- | |
| active | auto | | 0 | NO | NO | --- | |

PE3#show evpn mpls xconnect tunnel

| EVPN-MPLS Network tunnel Entries | | | | | | | |
|----------------------------------|-------------|-----------|----------|----------|---------------|---------|---------|
| Source | Destination | Status | Up/Down | Update | local-evpn-id | remote- | evpn-id |
| 10.10.10.3 | 10.10.10.2 | Installed | 00:13:15 | 00:04:12 | 1700 | 1800 | |
| 10.10.10.3 | 10.10.10.1 | Installed | 00:13:15 | 00:04:44 | 1700 | 1800 | |

Total number of entries are 2

PE4#show evpn load-balance all

| ESI | carving | weight | Revertive | AC-IF/PE AC-DF | PE-IP-ADDRESS Status | Redundancy | Service- |
|-------------------------------|---------|--------|-----------|-------------------|-------------------------|------------|----------|
| 00:12:22:33:00:00:00:55:66:77 | --- | | | | 10.10.10.3 | single- | |
| active | auto | | 0 | NO | NO | --- | |
| 00:12:22:33:00:00:00:55:66:77 | | | sa2.1 | | 10.10.10.4 | single- | |
| active | auto | | 0 | NO | NO | STANDBY | |

PE4#show evpn mpls xconnect tunnel

| EVPN-MPLS Network tunnel Entries | | | | | | | |
|----------------------------------|-------------|-----------|----------|----------|---------------|---------|---------|
| Source | Destination | Status | Up/Down | Update | local-evpn-id | remote- | evpn-id |
| 10.10.10.4 | 10.10.10.2 | Installed | 00:12:52 | 00:04:17 | 1700 | 1800 | |
| 10.10.10.4 | 10.10.10.1 | Installed | 00:12:52 | 00:04:17 | 1700 | 1800 | |

Total number of entries are 2

Single-Active ELAN

The following show output displays the active EVPN SR Tunnels and load balance for ELAN on PE1, PE2, PE3, and PE4 devices in the network [Figure 54](#) using the `show evpn mpls tunnel` and `show evpn load-balance all` commands.

```

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      evpn-id
=====
10.10.10.1      10.10.10.2      Installed   00:17:00     00:17:00    3000
10.10.10.1      10.10.10.4      Installed   00:18:10     00:18:10    3000
10.10.10.1      10.10.10.3      Installed   00:18:10     00:18:10    3000

Total number of entries are 3

PE1#show evpn load-balance all
ESI             AC-IF/PE      PE-IP-ADDRESS  Redundancy      Service-
carving  weight  Revertive  AC-DF  Status
=====
00:11:22:33:00:00:00:55:66:77  sa1.1      10.10.10.1      single-
active      auto      0      NO      NO      ACTIVE
00:11:22:33:00:00:00:55:66:77  ----      10.10.10.2      single-
active      auto      0      NO      NO      ----

PE2#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      evpn-id
=====
10.10.10.2      10.10.10.4      Installed   00:17:09     00:17:09    3000
10.10.10.2      10.10.10.3      Installed   00:17:09     00:17:09    3000
10.10.10.2      10.10.10.1      Installed   00:17:09     00:17:09    3000

Total number of entries are 3

PE2#show evpn load-balance all
ESI             AC-IF/PE      PE-IP-ADDRESS  Redundancy      Service-
carving  weight  Revertive  AC-DF  Status
=====
00:11:22:33:00:00:00:55:66:77  ----      10.10.10.1      single-
active      auto      0      NO      NO      ----
00:11:22:33:00:00:00:55:66:77  sa2.1      10.10.10.2      single-
active      auto      0      NO      NO      STANDBY

PE3#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      evpn-id
=====
10.10.10.3      10.10.10.2      Installed   00:17:11     00:17:11    3000
10.10.10.3      10.10.10.1      Installed   00:18:21     00:18:21    3000
10.10.10.3      10.10.10.4      Installed   00:29:15     00:28:54    3000

Total number of entries are 3

PE3#show evpn load-balance all
ESI             AC-IF/PE      PE-IP-ADDRESS  Redundancy      Service-
carving  weight  Revertive  AC-DF  Status
=====
00:12:22:33:00:00:00:55:66:77  sa1.1      10.10.10.3      single-
active      auto      0      NO      NO      ACTIVE
00:12:22:33:00:00:00:55:66:77  ----      10.10.10.4      single-
active      auto      0      NO      NO      ----

```

```
PE4#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      evpn-id
=====
10.10.10.4      10.10.10.2      Installed   00:17:13     00:17:13    3000
10.10.10.4      10.10.10.1      Installed   00:18:23     00:18:23    3000
10.10.10.4      10.10.10.3      Installed   00:29:18     00:29:14    3000
```

Total number of entries are 3

```
PE4#show evpn load-balance all
```

| ESI | carving | weight | Revertive | AC-IF/PE AC-DF | PE-IP-ADDRESS Status | Redundancy | Service- |
|-------------------------------|---------|--------|-----------|-------------------|-------------------------|------------|----------|
| 00:12:22:33:00:00:00:55:66:77 | active | auto | 0 | ---- | 10.10.10.3 | single- | ---- |
| 00:12:22:33:00:00:00:55:66:77 | active | auto | 0 | sa2.1 | 10.10.10.4 | single- | STANDBY |

Port-Active

The following show output displays the Ether Channel summary for CE1, CE2, PE1, PE2, PE3, and PE4 devices in the network [Figure 54](#) using the **show etherchannel summary** command.

```
CE1#show etherchannel summary
Aggregator po1 100001
Aggregator Type: Layer2
Admin Key: 0001 - Oper Key 0001
Link: xe48 (5049) sync: 0
Link: xe50 (5051) sync: 1
CE2#show etherchannel summary
Aggregator po1 100001
Aggregator Type: Layer2
Admin Key: 0001 - Oper Key 0001
Link: ge11 (5011) sync: 1
Link: xe23 (5023) sync: 0
PE1#show etherchannel summary
Aggregator po1 100001
Aggregator Type: Layer3
Admin Key: 0001 - Oper Key 0001
Link: xe0 (10004) sync: 0
PE2#show etherchannel summary
Aggregator po1 100001
Aggregator Type: Layer3
Admin Key: 0001 - Oper Key 0001
Link: xe8 (10029) sync: 1
PE3#show etherchannel summary
Aggregator po1 100001
Aggregator Type: Layer3
Admin Key: 0001 - Oper Key 0001
Link: xe2 (10003) sync: 0
PE4#show etherchannel summary
Aggregator po1 100001
Aggregator Type: Layer3
Admin Key: 0001 - Oper Key 0001
Link: xe11 (10012) sync: 1
```

The following show output displays the status of LDP sessions on PE1, PE2, PE3, PE4, P1, and P2 devices in the network [Figure 54](#) using the **show ldp session** command.

```
PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
```

t - TCP MSS not set/unset.
Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 10.10.10.2 | xe2 | Passive | OPERATIONAL | 30 | 00:06:57 |
| | 10.10.10.3 | xe14 | Passive | OPERATIONAL | 30 | 00:07:12 |
| | 10.10.10.4 | xe14 | Passive | OPERATIONAL | 30 | 00:06:42 |
| | 10.10.10.5 | xe14 | Passive | OPERATIONAL | 30 | 00:07:26 |
| | 10.10.10.6 | xe2 | Passive | OPERATIONAL | 30 | 00:06:36 |

PE2#show ldp session

Codes: m - MD5 password is not set/unset.
g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 10.10.10.1 | xe4 | Active | OPERATIONAL | 30 | 00:07:05 |
| | 10.10.10.3 | xe4 | Passive | OPERATIONAL | 30 | 00:07:05 |
| | 10.10.10.4 | xe4 | Passive | OPERATIONAL | 30 | 00:07:05 |
| | 10.10.10.5 | xe5 | Passive | OPERATIONAL | 30 | 00:07:03 |
| | 10.10.10.6 | xe4 | Passive | OPERATIONAL | 30 | 00:07:13 |

P1#show ldp session

Codes: m - MD5 password is not set/unset.
g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 10.10.10.1 | xe1 | Active | OPERATIONAL | 30 | 00:07:41 |
| | 10.10.10.2 | xe2 | Active | OPERATIONAL | 30 | 00:07:11 |
| | 10.10.10.3 | xe4 | Active | OPERATIONAL | 30 | 00:07:13 |
| | 10.10.10.4 | xe3 | Active | OPERATIONAL | 30 | 00:07:10 |

P2#show ldp session

Codes: m - MD5 password is not set/unset.
g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 10.10.10.1 | xe12 | Active | OPERATIONAL | 30 | 00:06:55 |
| | 10.10.10.2 | xe13 | Active | OPERATIONAL | 30 | 00:07:24 |
| | 10.10.10.3 | xe11 | Active | OPERATIONAL | 30 | 00:01:47 |
| | 10.10.10.4 | xe14 | Active | OPERATIONAL | 30 | 00:06:56 |

PE3#show ldp session

Codes: m - MD5 password is not set/unset.
g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 10.10.10.1 | xe5 | Active | OPERATIONAL | 30 | 00:07:35 |
| | 10.10.10.2 | xe5 | Active | OPERATIONAL | 30 | 00:07:20 |
| | 10.10.10.4 | xe5 | Passive | OPERATIONAL | 30 | 00:07:07 |
| | 10.10.10.5 | xe1 | Passive | OPERATIONAL | 30 | 00:07:21 |
| | 10.10.10.6 | xe5 | Passive | OPERATIONAL | 30 | 00:01:50 |

PE4#show ldp session

Codes: m - MD5 password is not set/unset.
g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 10.10.10.1 | xe0 | Active | OPERATIONAL | 30 | 00:07:09 |

| | | | | | |
|------------|-----|---------|-------------|----|----------|
| 10.10.10.2 | xe2 | Active | OPERATIONAL | 30 | 00:07:24 |
| 10.10.10.3 | xe0 | Active | OPERATIONAL | 30 | 00:07:11 |
| 10.10.10.5 | xe2 | Passive | OPERATIONAL | 30 | 00:07:22 |
| 10.10.10.6 | xe0 | Passive | OPERATIONAL | 30 | 00:07:03 |

The below show output displays the details about BGP L2VPN EVPN multihoming ES routes and Ethernet advertisement per ES for PE1, PE2, PE3, and PE4 devices in the network [Figure 54](#) using the **show bgp l2vpn evpn multihoming es-route** command.

```
PE1#show bgp l2vpn evpn multihoming es-route
```

```
RD[10.10.10.1:64512] VRF[evpn-gvrf-1]:
```

| ESI | DF | DP | weight | PE IP-Address | Encap | Peer IP | Algo | AC- |
|-------------------------------|----|----|--------|---------------|-------|------------|------|---------|
| 00:00:00:11:11:77:77:00:00:00 | | | | 10.10.10.1 | MPLS | ----- | DFT | no no 0 |
| 00:00:00:11:11:77:77:00:00:00 | | | | 10.10.10.2 | MPLS | 10.10.10.2 | DFT | no no 0 |

```
RD[10.10.10.2:64512]
```

| ESI | AC-DF | DP | weight | PE IP-Address | Encap | Peer IP | Algo |
|-------------------------------|-------|----|--------|---------------|-------|------------|------|
| 00:00:00:11:11:77:77:00:00:00 | | | | 10.10.10.2 | MPLS | 10.10.10.2 | DFT |
| | no | no | 0 | | | | |

```
PE2#show bgp l2vpn evpn multihoming es-route
```

```
RD[10.10.10.1:64512]
```

| ESI | PE IP-Address | Encap | Peer IP | Algo | AC-DF | DP | weight |
|-------------------------------|---------------|-------|------------|------|-------|----|--------|
| 00:00:00:11:11:77:77:00:00:00 | 10.10.10.1 | MPLS | 10.10.10.1 | DFT | no | no | 0 |

```
RD[10.10.10.2:64512] VRF[evpn-gvrf-1]:
```

| ESI | PE IP-Address | Encap | Peer IP | Algo | AC-DF | DP | weight |
|-------------------------------|---------------|-------|------------|------|-------|----|--------|
| 00:00:00:11:11:77:77:00:00:00 | 10.10.10.1 | MPLS | 10.10.10.1 | DFT | no | no | 0 |
| 00:00:00:11:11:77:77:00:00:00 | 10.10.10.2 | MPLS | ----- | DFT | no | no | 0 |

```
PE3#show bgp l2vpn evpn multihoming es-route
```

```
RD[10.10.10.3:64512] VRF[evpn-gvrf-1]:
```

| ESI | AC-DF | DP | weight | PE IP-Address | Encap | Peer IP | Algo |
|-------------------------------|-------|----|--------|---------------|-------|------------|------|
| 00:00:00:22:22:77:77:00:00:00 | | | | 10.10.10.3 | MPLS | ----- | DFT |
| | no | no | 0 | | | | |
| 00:00:00:22:22:77:77:00:00:00 | | | | 10.10.10.4 | MPLS | 10.10.10.4 | DFT |
| | no | no | 0 | | | | |

```
RD[10.10.10.4:64512]
```

| ESI | AC-DF | DP | weight | PE IP-Address | Encap | Peer IP | Algo |
|-------------------------------|-------|----|--------|---------------|-------|------------|------|
| 00:00:00:22:22:77:77:00:00:00 | | | | 10.10.10.4 | MPLS | 10.10.10.4 | DFT |
| | no | no | 0 | | | | |

```
PE4#show bgp l2vpn evpn multihoming es-route
```

```
RD[10.10.10.3:64512]
```

| ESI | PE IP-Address | Encap | Peer IP | Algo | AC-DF | DP | weight |
|-------------------------------|---------------|-------|------------|------|-------|----|--------|
| 00:00:00:22:22:77:77:00:00:00 | 10.10.10.3 | MPLS | 10.10.10.3 | DFT | no | no | 0 |

```
RD[10.10.10.4:64512] VRF[evpn-gvrf-1]:
```

| ESI | PE IP-Address | Encap | Peer IP | Algo | AC-DF | DP | weight |
|-------------------------------|---------------|-------|------------|------|-------|----|--------|
| 00:00:00:22:22:77:77:00:00:00 | 10.10.10.3 | MPLS | 10.10.10.3 | DFT | no | no | 0 |
| 00:00:00:22:22:77:77:00:00:00 | 10.10.10.4 | MPLS | ----- | DFT | no | no | 0 |

The following show output displays the details about Layer 2 Virtual Private Network (L2VPN) Ethernet Virtual Private Network (EVPN) routes on PE1, PE2, PE3, and PE4 devices in the network [Figure 54](#) using the **show bgp l2vpn evpn multihoming ethernet-ad-per-es** and **show bgp l2vpn evpn multihoming ethernet-ad-per-evi** commands.

```
PE1#show bgp l2vpn evpn multihoming ethernet-ad-per-es
```

```
RD[10.10.10.1:1700] VRF[vrf2]:
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:11:11:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.2 | MPLS | P flag |
| 00:00:00:22:22:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.3 | MPLS | B flag |
| 00:00:00:22:22:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.4 | MPLS | P flag |

```
RD[10.10.10.1:64512] VRF[evpn-gvrf-1]:
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:11:11:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.1 | MPLS | B flag |

```
RD[10.10.10.2:64512]
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:11:11:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.2 | MPLS | P flag |

```
RD[10.10.10.3:64512]
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:22:22:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.3 | MPLS | B flag |

```
RD[10.10.10.4:64512]
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:22:22:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.4 | MPLS | P flag |

```
PE1#show bgp l2vpn evpn multihoming ethernet-ad-per-evi
```

```
RD[10.10.10.1:1700] VRF[vrf2]:
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|---------|------------|------------|-------|--------|
| 00:00:00:11:11:77:77:00:00:00 | 1800 | 27522 | 10.10.10.1 | MPLS | B flag |
| 00:00:00:11:11:77:77:00:00:00 | 1800 | 27520 | 10.10.10.2 | MPLS | P flag |
| 00:00:00:22:22:77:77:00:00:00 | 1700 | 27520 | 10.10.10.3 | MPLS | B flag |
| 00:00:00:22:22:77:77:00:00:00 | 1700 | 27520 | 10.10.10.4 | MPLS | P flag |

```
RD[10.10.10.2:1700]
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|---------|------------|------------|-------|--------|
| 00:00:00:11:11:77:77:00:00:00 | 1800 | 27520 | 10.10.10.2 | MPLS | P flag |

```
RD[10.10.10.3:1700]
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|---------|------------|------------|-------|--------|
| 00:00:00:22:22:77:77:00:00:00 | 1700 | 27520 | 10.10.10.3 | MPLS | B flag |

```
RD[10.10.10.4:1700]
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|---------|------------|------------|-------|--------|
| 00:00:00:22:22:77:77:00:00:00 | 1700 | 27520 | 10.10.10.4 | MPLS | P flag |

```
PE2#show bgp l2vpn evpn multihoming ethernet-ad-per-es
```

```
RD[10.10.10.1:64512]
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:11:11:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.1 | MPLS | B flag |

```
RD[10.10.10.2:1700] VRF[vrf2]:
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:11:11:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.1 | MPLS | B flag |
| 00:00:00:22:22:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.3 | MPLS | B flag |
| 00:00:00:22:22:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.4 | MPLS | P flag |

```
RD[10.10.10.2:64512] VRF[evpn-gvrf-1]:
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:11:11:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.2 | MPLS | P flag |

```
RD[10.10.10.3:64512]
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:22:22:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.3 | MPLS | B flag |

```
RD[10.10.10.4:64512]
```

| ESI | Eth-Tag | VNID/LABEL | Nexthop IP | Encap | Flags |
|-------------------------------|------------|------------|------------|-------|--------|
| 00:00:00:22:22:77:77:00:00:00 | 4294967295 | 440336 | 10.10.10.4 | MPLS | P flag |

```
PE2#show bgp l2vpn evpn multihoming ethernet-ad-per-evi
```

```
RD[10.10.10.1:1700]
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:11:11:77:77:00:00:00  1800    27522      10.10.10.1  MPLS   B flag

RD[10.10.10.2:1700] VRF[vrf2]:
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:11:11:77:77:00:00:00  1800    27522      10.10.10.1  MPLS   B flag
00:00:00:11:11:77:77:00:00:00  1800    27520      10.10.10.2  MPLS   P flag
00:00:00:22:22:77:77:00:00:00  1700    27520      10.10.10.3  MPLS   B flag
00:00:00:22:22:77:77:00:00:00  1700    27520      10.10.10.4  MPLS   P flag

RD[10.10.10.3:1700]
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:22:22:77:77:00:00:00  1700    27520      10.10.10.3  MPLS   B flag

RD[10.10.10.4:1700]
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:22:22:77:77:00:00:00  1700    27520      10.10.10.4  MPLS   P flag
```

```
PE3#show bgp l2vpn evpn multihoming ethernet-ad-per-es
```

```
RD[10.10.10.1:64512]
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:11:11:77:77:00:00:00  4294967295  440336    10.10.10.1  MPLS   B flag

RD[10.10.10.2:64512]
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:11:11:77:77:00:00:00  4294967295  440336    10.10.10.2  MPLS   P flag

RD[10.10.10.3:1700] VRF[vrf2]:
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:11:11:77:77:00:00:00  4294967295  440336    10.10.10.2  MPLS   P flag
00:00:00:11:11:77:77:00:00:00  4294967295  440336    10.10.10.1  MPLS   B flag
00:00:00:22:22:77:77:00:00:00  4294967295  440336    10.10.10.4  MPLS   P flag

RD[10.10.10.3:64512] VRF[evpn-gvrf-1]:
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:22:22:77:77:00:00:00  4294967295  440336    10.10.10.3  MPLS   B flag

RD[10.10.10.4:64512]
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:22:22:77:77:00:00:00  4294967295  440336    10.10.10.4  MPLS   P flag
```

```
PE3#show bgp l2vpn evpn multihoming ethernet-ad-per-evi
```

```
RD[10.10.10.1:1700]
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:11:11:77:77:00:00:00  1800    27522      10.10.10.1  MPLS   B flag

RD[10.10.10.2:1700]
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:11:11:77:77:00:00:00  1800    27520      10.10.10.2  MPLS   P flag

RD[10.10.10.3:1700] VRF[vrf2]:
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:11:11:77:77:00:00:00  1800    27520      10.10.10.2  MPLS   P flag
00:00:00:11:11:77:77:00:00:00  1800    27522      10.10.10.1  MPLS   B flag
00:00:00:22:22:77:77:00:00:00  1700    27520      10.10.10.3  MPLS   B flag
00:00:00:22:22:77:77:00:00:00  1700    27520      10.10.10.4  MPLS   P flag

RD[10.10.10.4:1700]
ESI                               Eth-Tag  VNID/LABEL  Nexthop IP  Encap  Flags
00:00:00:22:22:77:77:00:00:00  1700    27520      10.10.10.4  MPLS   P flag
```

```
PE4#show bgp l2vpn evpn multihoming ethernet-ad-per-es
```

```

RD[10.10.10.1:64512]
ESI                               Eth-Tag   VNID/LABEL   Nexthop IP   Encap   Flags
00:00:00:11:11:77:77:00:00:00  4294967295  440336      10.10.10.1   MPLS    B flag

RD[10.10.10.2:64512]
ESI                               Eth-Tag   VNID/LABEL   Nexthop IP   Encap   Flags
00:00:00:11:11:77:77:00:00:00  4294967295  440336      10.10.10.2   MPLS    P flag

RD[10.10.10.3:64512]
ESI                               Eth-Tag   VNID/LABEL   Nexthop IP   Encap   Flags
00:00:00:22:22:77:77:00:00:00  4294967295  440336      10.10.10.3   MPLS    B flag

RD[10.10.10.4:1700] VRF[vrf2]:
ESI                               Eth-Tag   VNID/LABEL   Nexthop IP   Encap   Flags
00:00:00:11:11:77:77:00:00:00  4294967295  440336      10.10.10.1   MPLS    B flag
00:00:00:11:11:77:77:00:00:00  4294967295  440336      10.10.10.2   MPLS    P flag
00:00:00:22:22:77:77:00:00:00  4294967295  440336      10.10.10.3   MPLS    B flag

RD[10.10.10.4:64512] VRF[evpn-gvrf-1]:
ESI                               Eth-Tag   VNID/LABEL   Nexthop IP   Encap   Flags
00:00:00:22:22:77:77:00:00:00  4294967295  440336      10.10.10.4   MPLS    P flag

PE4#show bgp l2vpn evpn multihoming ethernet-ad-per-evi

RD[10.10.10.1:1700]
ESI                               Eth-Tag   VNID/LABEL   Nexthop IP   Encap   Flags
00:00:00:11:11:77:77:00:00:00  1800      27522      10.10.10.1   MPLS    B flag

RD[10.10.10.2:1700]
ESI                               Eth-Tag   VNID/LABEL   Nexthop IP   Encap   Flags
00:00:00:11:11:77:77:00:00:00  1800      27520      10.10.10.2   MPLS    P flag

RD[10.10.10.3:1700]
ESI                               Eth-Tag   VNID/LABEL   Nexthop IP   Encap   Flags
00:00:00:22:22:77:77:00:00:00  1700      27520      10.10.10.3   MPLS    B flag

RD[10.10.10.4:1700] VRF[vrf2]:
ESI                               Eth-Tag   VNID/LABEL   Nexthop IP   Encap   Flags
00:00:00:11:11:77:77:00:00:00  1800      27522      10.10.10.1   MPLS    B flag
00:00:00:11:11:77:77:00:00:00  1800      27520      10.10.10.2   MPLS    P flag
00:00:00:22:22:77:77:00:00:00  1700      27520      10.10.10.3   MPLS    B flag
00:00:00:22:22:77:77:00:00:00  1700      27520      10.10.10.4   MPLS    P flag

```

Port-Active ELINE

The following show output displays the active EVPN MPLS Tunnels for ELINE on PE1, PE2, PE3, and PE4 devices in the network [Figure 54](#) using the `show evpn mpls xconnect tunnel` command.

```

PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====
===
10.10.10.1  10.10.10.3      AC-Down     00:31:41     00:31:41     4             1700
10.10.10.1  10.10.10.4      AC-Down     00:31:41     00:31:41     4             1700
Total number of entries are 2

PE2#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====
===
10.10.10.2  10.10.10.3      Installed   00:12:21     00:11:40     1800          1700
10.10.10.2  10.10.10.4      Installed   00:17:43     00:17:37     1800          1700

```

Total number of entries are 2

PE3#show evpn mpls xconnect tunnel

EVPN-MPLS Network tunnel Entries

| Source evpn-id | Destination | Status | Up/Down | Update | local-evpn-id | remote- |
|-------------------|-------------|---------|----------|----------|---------------|---------|
| 10.10.10.3 | 10.10.10.1 | AC-Down | 00:12:26 | 00:12:26 | 1700 | 1800 |
| 10.10.10.3 | 10.10.10.2 | AC-Down | 00:12:26 | 00:12:26 | 1700 | 1800 |

Total number of entries are 2

PE4#show evpn mpls xconnect tunnel

EVPN-MPLS Network tunnel Entries

| Source evpn-id | Destination | Status | Up/Down | Update | local-evpn-id | remote- |
|-------------------|-------------|-----------|----------|----------|---------------|---------|
| 10.10.10.4 | 10.10.10.1 | Installed | 00:12:28 | 00:12:28 | 1700 | 1800 |
| 10.10.10.4 | 10.10.10.2 | Installed | 00:12:28 | 00:12:28 | 1700 | 1800 |

Total number of entries are 2

Port-Active ELAN

The following show outputs provide validation for ELAN configurations.

The following show output displays the active EVPN MPLS Tunnels for ELAN on PE1, PE2, PE3, and PE4 devices in the network [Figure 54](#) using the show evpn mpls tunnel command.

PE1#show evpn mpls tunnel

EVPN-MPLS Network tunnel Entries

| Source | Destination | Status | Up/Down | Update | evpn-id |
|------------|-------------|-----------|----------|----------|---------|
| 10.10.10.1 | 10.10.10.4 | Installed | 00:02:35 | 00:02:35 | 3000 |
| 10.10.10.1 | 10.10.10.3 | Installed | 00:03:00 | 00:03:00 | 3000 |
| 10.10.10.1 | 10.10.10.2 | Installed | 00:03:26 | 00:03:26 | 3000 |

Total number of entries are 3

PE2#show evpn mpls tunnel

EVPN-MPLS Network tunnel Entries

| Source | Destination | Status | Up/Down | Update | evpn-id |
|------------|-------------|-----------|----------|----------|---------|
| 10.10.10.2 | 10.10.10.4 | Installed | 00:02:45 | 00:02:45 | 3000 |
| 10.10.10.2 | 10.10.10.3 | Installed | 00:03:10 | 00:03:10 | 3000 |
| 10.10.10.2 | 10.10.10.1 | Installed | 00:03:36 | 00:03:36 | 3000 |

Total number of entries are 3

PE3#show evpn mpls tunnel

EVPN-MPLS Network tunnel Entries

| Source | Destination | Status | Up/Down | Update | evpn-id |
|------------|-------------|-----------|----------|----------|---------|
| 10.10.10.3 | 10.10.10.4 | Installed | 00:02:56 | 00:02:56 | 3000 |
| 10.10.10.3 | 10.10.10.2 | Installed | 00:03:22 | 00:03:22 | 3000 |


```

10.10.10.3      10.10.10.1      Installed      00:03:22      00:03:22      3000

Total number of entries are 3

PE4#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status          Up/Down          Update          evpn-id
=====
10.10.10.4      10.10.10.3      Installed      00:03:00      00:03:00      3000
10.10.10.4      10.10.10.1      Installed      00:03:00      00:03:00      3000
10.10.10.4      10.10.10.2      Installed      00:03:00      00:03:00      3000

Total number of entries are 3

```

The following show output displays the EVPN active multi-homed and load-balanced details on PE1, PE2, PE3, and PE4 devices in the network [Figure 54](#) using the show evpn load-balance port-active and show evpn multi-homing all commands.

```

PE1#show evpn load-balance port-active
ESI          AC-IF/PE      PE-IP-ADDRESS      Redundancy      Service-
carving      weight      Revertive      AC-DF      Status
=====
00:00:00:11:11:77:77:00:00:00  LOCAL      10.10.10.1      port-
active      auto      0      NO      NA      STANDBY
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.2      port-
active      auto      0      NO      NA      ACTIVE
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.3      port-active      ----      --
--      ----      ----      STANDBY
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.4      port-active      ----      --
--      ----      ----      ACTIVE

PE1#show evpn multi-homing all
ESI          Access-IF      PE-IP-ADDRESS
=====
00:00:00:11:11:77:77:00:00:00  po1      10.10.10.1
00:00:00:11:11:77:77:00:00:00  ----      10.10.10.2
00:00:00:22:22:77:77:00:00:00  ----      10.10.10.3
00:00:00:22:22:77:77:00:00:00  ----      10.10.10.4
Total number of entries are 4

PE2#show evpn load-balance port-active
ESI          AC-IF/PE      PE-IP-ADDRESS      Redundancy      Service-
carving      weight      Revertive      AC-DF      Status
=====
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.1      port-
active      auto      0      NO      NA      STANDBY
00:00:00:11:11:77:77:00:00:00  LOCAL      10.10.10.2      port-
active      auto      0      NO      NA      ACTIVE
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.3      port-active      ---      ----
----      ----      STANDBY
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.4      port-active      ----      ----
----      ----      ACTIVE

PE2#show evpn multi-homing all
ESI          Access-IF      PE-IP-ADDRESS
=====
00:00:00:11:11:77:77:00:00:00  ----      10.10.10.1
00:00:00:11:11:77:77:00:00:00  po1      10.10.10.2
00:00:00:22:22:77:77:00:00:00  ----      10.10.10.3
00:00:00:22:22:77:77:00:00:00  ----      10.10.10.4
Total number of entries are 4

```

```

PE3#show evpn load-balance port-active
ESI              AC-IF/PE      PE-IP-ADDRESS    Redundancy      Service-
carving  weight  Revertive  AC-DF  Status
=====
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.1      port-active      ----      --
--              ----      STANDBY
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.2      port-active      ----      --
--              ----      ACTIVE
00:00:00:22:22:77:77:00:00:00  LOCAL      10.10.10.3      port-
active          auto          0          NO          NA      STANDBY
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.4      port-
active          auto          0          NO          NA      ACTIVE

PE3#show evpn multi-homing all
ESI              Access-IF      PE-IP-ADDRESS
=====
00:00:00:11:11:77:77:00:00:00  ----      10.10.10.1
00:00:00:11:11:77:77:00:00:00  ----      10.10.10.2
00:00:00:22:22:77:77:00:00:00  po1       10.10.10.3
00:00:00:22:22:77:77:00:00:00  ----      10.10.10.4
Total number of entries are 4

PE4#show evpn load-balance port-active
ESI              AC-IF/PE      PE-IP-ADDRESS    Redundancy      Service-
carving  weight  Revertive  AC-DF  Status
=====
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.1      port-active      ----      --
--              ----      STANDBY
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.2      port-active      ----      --
--              ----      ACTIVE
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.3      port-
active          auto          0          NO          NA      STANDBY
00:00:00:22:22:77:77:00:00:00  LOCAL      10.10.10.4      port-
active          auto          0          NO          NA      ACTIVE

PE4#show evpn multi-homing all
ESI              Access-IF      PE-IP-ADDRESS
=====
00:00:00:11:11:77:77:00:00:00  ----      10.10.10.1
00:00:00:11:11:77:77:00:00:00  ----      10.10.10.2
00:00:00:22:22:77:77:00:00:00  ----      10.10.10.3
00:00:00:22:22:77:77:00:00:00  po1       10.10.10.4

```


EVPN SR Active-Standby Multi-Homing Configuration

EVPN SR Active-Standby Multi-Homing Configuration

This section illustrates the Multi-Homed setup for the EVPN Segment Routing (SR) Active-Standby configuration, showcasing examples for both ELINE and ELAN services with SR as the underlay MPLS path.

EVPN SR Active-Standby MH Topology

The below topology consists of customer edge routers CE1 and CE2, along with IPv4 Provider Edge routers PE1, PE2, PE3, and PE4, all interconnected through the core routers P1 and P2 in the IPv4 MPLS provider network.

Figure 55. EVPN MPLS AS MH Configuration

Customer Edge (CE) Routers Configurations

CE1

The following configuration steps under CE1 are set up to enable VLANs and configure interfaces for carrying VLAN traffic.

| | |
|--|--|
| CE1#configure terminal | Enter configure mode. |
| CE1(config)#bridge 1 protocol ieee vlan-bridge | Set up bridge 1 to use the IEEE VLAN bridge protocol. |
| CE1(config)#vlan 2-100 bridge 1 state enable | Configure VLANs from 2-100 and associate them with bridge 1. |
| CE1(config)#interface xe54 | Enter interface mode xe54. |
| CE1(config-if)#switchport | Configure the interface xe54 as a Layer 2 switch port. |
| CE1(config-if)#bridge-group 1 | Associate xe54 to bridge 1. |
| CE1(config-if)#switchport mode trunk | Configure xe54 as a trunk port. |
| CE1(config-if)#switchport trunk allowed vlan all | Allow all configured VLANs on the trunk interface xe54. |
| CE1(config-if)#exit | Exit interface mode xe54. |
| CE1(config)#interface po1 | Enter interface mode and configure LAG interface port-channel 1 (po1). |
| CE1(config-if)#switchport | Configures port-channel 1 as a Layer 2 switch port. |
| CE1(config-if)#bridge-group 1 | Associate po1 to bridge 1. |
| CE1(config-if)#switchport mode trunk | Configure po1 as a trunk port. |
| CE1(config-if)#switchport trunk allowed vlan all | Allow all configured VLANs on the trunk port-channel po1. |
| CE1(config-if)#exit | Exit interface mode po1. |

| | |
|--|---|
| CE1(config)#interface xe48 | Enter interface mode <code>xe48</code> . |
| CE1(config-if)#lACP timeout short | Configure LACP timeout as <code>short</code> . |
| CE1(config-if)#channel-group 1 mode active | Add member to the LAG interface. |
| CE1(config-if)#exit | Exit interface mode <code>xe48</code> . |
| CE1(config-if)#interface xe50 | Enter interface mode <code>xe50</code> . |
| CE1(config-if)#lACP timeout short | Configure LACP timeout as <code>short</code> . |
| CE1(config-if)#channel-group 1 mode active | Add member to the LAG interface. |
| CE1(config-if)#commit | Commit the transaction. |
| CE1(config-if)#end | Exit interface mode <code>xe50</code> and configure mode. |

CE2

The following configuration steps under CE2 are set up to enable VLANs and configure interfaces for carrying VLAN traffic.

| | |
|--|---|
| CE2#configure terminal | Enter configure mode. |
| CE2(config)#bridge 1 protocol ieee vlan-bridge | Set up bridge 1 to use the IEEE VLAN bridge protocol. |
| CE2(config)#vlan 2-100 bridge 1 state enable | Configure VLANs from 2-100 and associate them with bridge 1. |
| CE2(config)#interface xe24 | Enter interface mode <code>xe24</code> . |
| CE2(config-if)#switchport | Configure the interface <code>xe24</code> as a Layer 2 switch port. |
| CE2(config-if)#bridge-group 1 | Associate <code>xe24</code> to bridge 1. |
| CE2(config-if)#switchport mode trunk | Configure <code>xe24</code> as a trunk port. |
| CE2(config-if)#switchport trunk allowed vlan all | Allow all configured VLANs on the trunk interface <code>xe24</code> . |
| CE2(config-if)#exit | Exit interface mode <code>xe24</code> . |
| CE2(config)#interface po1 | Enter interface mode and configure LAG interface port-channel 1 (<code>po1</code>). |
| CE2(config-if)#switchport | Configures port-channel 1 as a Layer 2 switch port. |
| CE2(config-if)#bridge-group 1 | Associate <code>po1</code> to bridge 1. |
| CE2(config-if)#switchport mode trunk | Configure <code>po1</code> as a trunk port. |
| CE2(config-if)#switchport trunk allowed vlan all | Allow all configured VLANs on the trunk port-channel <code>po1</code> . |
| CE2(config-if)#exit | Exit interface mode <code>po1</code> . |
| CE2(config)#interface xe22 | Enter interface mode <code>xe22</code> . |
| CE2(config-if)#lACP timeout short | Configure LACP timeout as <code>short</code> . |
| CE2(config-if)#channel-group 1 mode active | Add member to the LAG interface. |
| CE2(config-if)#exit | Exit interface mode <code>xe22</code> . |

| | |
|--|---|
| CE2(config-if)#interface xe23 | Enter interface mode <code>xe23</code> . |
| CE2(config-if)#lACP timeout short | Configure LACP timeout as <code>short</code> . |
| CE2(config-if)#channel-group 1 mode active | Add member to the LAG interface. |
| CE2(config-if)#commit | Commit the transaction. |
| CE2(config-if)#end | Exit interface mode <code>xe23</code> and configure mode. |

Provider Edge Routers Configurations

PE1

PE1: Loopback Interface

The configuration on PE1 for a loopback interface with IP address `10.10.10.1/32` secondary is set up to provide IP connectivity for the ISIS router.

| | |
|---|---|
| PE1#configure terminal | Enter configure mode. |
| PE1(config)#interface lo | Enter the interface mode for the loopback interface <code>lo</code> . |
| PE1(config-if)#ip address 10.10.10.1/32 secondary | Configure a secondary IP address, <code>10.10.10.1/32</code> , on the loopback interface. |
| PE1(config-if)#ip router isis 1 | Enable ISIS routing on a loopback interface <code>lo</code> for area 1. |
| PE1(config-if)#prefix-sid index 800 | Configure a prefix segment identifier (prefix-SID) index value as <code>800</code> . |
| PE1(config-if)#exit | Exit interface mode <code>lo</code> . |
| PE1(config)#commit | Commit the transaction. |

PE1: Configure SR

The following configurations aim to activate Segment Routing (SR) on PE1 and make MPLS the preferred method for segment routing, optimizing routing efficiency.

| | |
|-------------------------------|---|
| PE1(config)#segment-routing | Configure segment routing on PE1 device. |
| PE1(config-sr)#mpls sr-prefer | Set MPLS as the preferred segment routing protocol over others. |
| PE1(config-sr)#exit | Exit the router SR mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: Global LDP

The configuration on PE1 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|------------------------|----------------------------|
| PE1(config)#router ldp | Enter the Router LDP mode. |
|------------------------|----------------------------|

| | |
|--|---|
| PE1(config-router)#router-id 10.10.10.1 | Set the router ID for LDP to 10.10.10.1. |
| PE1(config-router)#transport-address ipv4 10.10.10.1 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| PE1(config-router)#targeted-peer ipv4 10.10.10.2 | Configure targeted peer for LDP using IPv4 addresses. |
| PE1(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE1(config-router)#targeted-peer ipv4 10.10.10.3 | Configure targeted peer for LDP using IPv4 addresses. |
| PE1(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE1(config-router)#targeted-peer ipv4 10.10.10.4 | Configure targeted peer for LDP using IPv4 addresses. |
| PE1(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE1(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: Global EVPN MPLS Command

The configuration on PE1 for the Global EVPN MPLS, includes activating EVPN MPLS, defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multi-homing, and activating EVPN MPLS multi-homing functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|---|--|
| PE1(config)#evpn mpls enable | Activate the EVPN MPLS functionality on PE1, enabling it to participate in EVPN MPLS services. |
| PE1(config)#commit | Commit candidate configuration to be running configuration. |
| PE1(config)#evpn mpls vtep-ip-global 10.10.10.1 | Configure the global VTEP IP address 10.10.10.1, associating it with the loopback IP. |
| PE1(config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter for EVPN MPLS multi-homing. |
| PE1(config)#evpn mpls multihoming enable | Activate the EVPN MPLS multi-homing functionality, allowing PE1 to support multi-homed EVPN MPLS services. |
| PE1(config)#commit | Commit the transaction. |

PE1: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE1 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|---------------------------------------|--|
| PE1(config)#interface xe1 | Enter interface mode xe1. |
| PE1(config-if)#ip address 10.1.1.1/30 | Configure an IP address, 10.1.1.1/30, on the interface xe1. |
| PE1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE1(config-if)#ip router isis 1 | Enable ISIS IPv4 routing on an interface xe1. |
| PE1(config-if)#exit | Exit interface mode xe1. |
| PE1(config)#commit | Commit the transaction. |
| PE1(config)#interface xe2 | Enter interface mode xe2. |
| PE1(config-if)#ip address 10.1.2.1/30 | Configure an IP address, 10.1.2.1/30, on the interface xe2. |
| PE1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE1(config-if)#ip router isis 1 | Enable ISIS IPv4 routing on an interface xe2. |
| PE1(config-if)#exit | Exit interface mode xe2. |
| PE1(config)#commit | Commit the transaction. |

PE1: ISIS Configuration

The below configuration is performed to set up ISIS on PE1, to enable MPLS Traffic Engineering, Segment Routing, and other related features for efficient routing and network management.

| | |
|--|--|
| PE1(config)#router isis 1 | Enter router ISIS mode. |
| PE1(config-router)#is-type level-1-2 | Configure IS-Type as Level-1-2 specifies that the router will participate in both Level-1 and Level-2 areas within the ISIS network. |
| PE1(config-router)#metric-style wide | Configure the new style of metric type as wide. |
| PE1(config-router)#mpls traffic-eng router-id 10.10.10.1 | Configure the router's MPLS Traffic Engineering (TE) router ID TLV to 10.10.10.1, which is used for MPLS-TE path calculations. |
| PE1(config-router)#mpls traffic-eng level-1 | Enable MPLS-TE for IS-Type Level-1 routing. |
| PE1(config-router)#mpls traffic-eng level-2 | Enable MPLS-TE for IS-Type Level-2 routing. |
| PE1(config-router)#capability cspf | Enable Constraint Shortest Path First (CSPF) computation for traffic engineering. |
| PE1(config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |

| | |
|---|--|
| PE1(config-router)#fast-reroute ti-lfa level-1 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| PE1(config-router)#fast-reroute ti-lfa level-2 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| PE1(config-router)#bfd all-interfaces | Configure BFD on all interfaces for fast link failure detection. |
| PE1(config-router)#net 49.0000.0000.0001.00 | Set a Network Entity Title (NET) for this ISIS instance, specifying the area address and the system ID. |
| PE1(config-router)#isis segment-routing global block 17000 23500 | Enable ISIS SR globally and allocates label blocks for Segment Routing. |
| PE1(config-router)#segment-routing mpls | Enable SR ISIS. |
| PE1(config-router)#exit | Exit router ISIS mode and return to configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: BGP Configuration

The below BGP configuration on PE1 is established to enable BGP routing with ASN 65010, set the BGP router ID, define iBGP neighbors, configure BFD, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|---|--|
| PE1(config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| PE1(config-router)#bgp router-id 10.10.10.1 | Configure BGP router ID 10.10.10.1, identifying PE1 within the BGP network. |
| PE1(config-router)#neighbor 10.10.10.2 remote-as 65010 | Configure neighbor 10.10.10.2 as an iBGP neighbor with their remote AS number 65010. |
| PE1(config-router)#neighbor 10.10.10.2 update-source lo | Configure neighbor 10.10.10.2 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE1(config-router)#neighbor 10.10.10.3 remote-as 65010 | Configure neighbor 10.10.10.3 as an iBGP neighbor with their remote AS number 65010. |
| PE1(config-router)#neighbor 10.10.10.3 update-source lo | Configure neighbor 10.10.10.3 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE1(config-router)#neighbor 10.10.10.4 remote-as 65010 | Configure neighbor 10.10.10.4 as an iBGP neighbor with their remote AS number 65010. |
| PE1(config-router)#neighbor 10.10.10.4 update-source lo | Configure neighbor 10.10.10.4 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |

| | |
|---|---|
| PE1(config-router)#neighbor 10.10.10.2 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE1(config-router)#neighbor 10.10.10.3 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE1(config-router)#neighbor 10.10.10.4 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE1(config-router)#neighbor 10.10.10.2 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE1(config-router)#neighbor 10.10.10.3 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE1(config-router)#neighbor 10.10.10.4 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE1(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE1(config-router-af)#neighbor 10.10.10.2 activate | Activate EVPN for iBGP neighbor 10.10.10.2 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE1(config-router-af)#neighbor 10.10.10.3 activate | Activate EVPN for iBGP neighbor 10.10.10.3 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE1(config-router-af)#neighbor 10.10.10.4 activate | Activate EVPN for iBGP neighbor 10.10.10.4 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE1(config-router-af)#exit | Exit address family mode and return to the router BGP mode. |
| PE1(config-router)#commit | Commit the transaction. |
| PE1(config-router)#exit | Exit router BGP mode and return to the configure mode. |

PE1: MAC VRF Configuration

The below MAC VRF configuration on PE1 is carried out to define and set up VRFs named `vrf2` and `vpls1001` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|---|--|
| PE1(config)#mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE1(config-vrf)#rd 10.10.10.1:1700 | Configure Route-Distinguisher value of 10.10.10.1:1700. |
| PE1(config-vrf)#route-target both 1700:1700 | Configure import and export values for the <code>vrf2</code> as 1700:1700. |
| PE1(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE1(config)#mac vrf vpls1001 | Enter VRF mode named <code>vpls1001</code> . |
| PE1(config-vrf)#rd 10.10.10.1:1001 | Configure Route-Distinguisher value of 10.10.10.1:1001. |

| | |
|---|---|
| PE1(config-vrf)#route-target both 1001:1001 | Configure import and export values for the vpls1001 as 1001:1001. |
| PE1(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: EVPN and VRF Mapping

The below EVPN and VRF mapping configuration on PE1 is performed to establish mappings between EVPN identifiers and VRFs, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|--|--|
| PE1(config)#evpn mpls id 1800 xconnect target-mpls-id 1700 | Configure the EVPN-VPWS identifier with a source identifier of 1800 and a target identifier of 1700. |
| PE1(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Map VRF vrf2 to the EVPN-VPWS identifier |
| PE1(config-evpn-mpls)#commit | Commit the transaction. |
| PE1(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| PE1(config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier an identifier of 3000. |
| PE1(config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Map VRF vpls1001 to the EVPN-VPWS identifier |
| PE1(config-evpn-mpls)#commit | Commit the transaction. |
| PE1(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE1: Access Port Configuration for Port-active

The below access port configuration for port-active mode on PE1 is carried out to configure various parameters including system-mac, load balancing, service carving preferences, and EVPN settings for efficient network access and connectivity.

| | |
|--|--|
| PE1(config)#interface po1 | Enter the port channel interface mode for po1 |
| PE1(config-if)#load-interval 30 | Set the load interval to 30. |
| PE1(config-if)#evpn multi-homed system-mac 0000.1111.7777 load-balancing port-active | Configure the system-mac address 0000.1111.7777 for port-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE1(config-if-es)#service-carving auto | Configure service carving as auto. |
| PE1(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE1(config-if)#exit | Exit interface mode po1 and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

| | |
|--|---|
| PE1(config)#interface po1.1 switchport | Create a Layer 2 sub-interface po1.1 within the port channel. |
| PE1(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE1(config-if)#load-interval 30 | Set the load interval to 30. |
| PE1(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE1(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE1(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE1(config-if)#exit | Exit interface mode po1.1 and return to the configure mode. |
| PE1(config)#interface xe0 | Enter the interface mode for xe0. |
| PE1(config-if)#speed 10g | Set the speed to 10g. |
| PE1(config-if)#channel-group 1 mode active | Attach LAG interface po1. |
| PE1(config-if)#exit | Exit interface mode xe0 and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: Access Port Configuration for Single-active

The below access port configuration for single-active mode on PE1 is implemented to set up various parameters, including Ethernet Segment Identifier (ESI) settings, service carving preferences, and EVPN configurations, ensuring efficient network access and connectivity.

| | |
|---|--|
| PE1(config)#interface sa1 | Enter the single active interface mode for sa1 |
| PE1(config-if)#load-interval 30 | Set the load interval to 30. |
| PE1(config-if)#evpn multi-homed esi 00:00:11:11:77:77 load-balancing single-active | Configure the ESI with the value with the value 00:00:11:11:77:77 for single-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE1(config-if-es)#service-carving auto | Configure service carving as auto, allowing automatic determination of service distribution preferences. |
| PE1(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE1(config-if)#exit | Exit interface mode sa1 and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |
| PE1(config)#interface sa1.1 switchport | Create a Layer 2 sub-interface sa1.1 within the port channel. |
| PE1(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE1(config-if)#load-interval 30 | Set the load interval to 30. |

| | |
|---|---|
| PE1(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE1(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE1(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE1(config-if)#exit | Exit interface mode <code>sa1.1</code> and return to the configure mode. |
| PE1(config)#interface xe0 | Enter the interface mode for <code>xe0</code> . |
| PE1(config-if)#speed 10g | Set the speed to 10g. |
| PE1(config-if)#static-channel-group 1 | Attach the static-channel-group 1, the LAG interface <code>sa1</code> to <code>xe0</code> . |
| PE1(config-if)#exit | Exit interface mode <code>xe0</code> and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE2

PE2: Loopback Interface

The configuration on PE2 for a loopback interface with IP address `10.10.10.2/32` secondary is set up to provide IP connectivity for the ISIS router.

| | |
|---|---|
| PE2#configure terminal | Enter configure mode. |
| PE2(config)#interface lo | Enter the interface mode for the loopback interface <code>lo</code> . |
| PE2(config-if)#ip address 10.10.10.2/32 secondary | Configure a secondary IP address, <code>10.10.10.2/32</code> , on the loopback interface. |
| PE2(config-if)#ip router isis 1 | Enable ISIS routing on a loopback interface <code>lo</code> for area 1. |
| PE2(config-if)#prefix-sid index 800 | Configure a prefix segment identifier (prefix-SID) index value as 800. |
| PE2(config-if)#exit | Exit interface mode <code>lo</code> . |
| PE2(config)#commit | Commit the transaction. |

PE2: Configure SR

The following configurations aim to activate Segment Routing (SR) on PE2 and make MPLS the preferred method for segment routing, optimizing routing efficiency.

| | |
|-------------------------------|---|
| PE2(config)#segment-routing | Configure segment routing on PE2 device. |
| PE2(config-sr)#mpls sr-prefer | Set MPLS as the preferred segment routing protocol over others. |

| | |
|---------------------|--------------------------|
| PE2(config-sr)#exit | Exit the router SR mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: Global LDP

The configuration on PE2 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|--|---|
| PE2(config)#router ldp | Enter the Router LDP mode. |
| PE2(config-router)#router-id 10.10.10.2 | Set the router ID for LDP to 10.10.10.2. |
| PE2(config-router)#transport-address ipv4 10.10.10.2 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| PE2(config-router)#targeted-peer ipv4 10.10.10.1 | Configure targeted peer for LDP using IPv4 addresses. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE2(config-router)#targeted-peer ipv4 10.10.10.3 | Configure targeted peer for LDP using IPv4 addresses. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE2(config-router)#targeted-peer ipv4 10.10.10.4 | Configure targeted peer for LDP using IPv4 addresses. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE2(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: Global EVPN MPLS Command

The configuration on PE2 for the Global EVPN MPLS, includes activating EVPN MPLS, defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multi-homing, and activating EVPN MPLS multi-homing functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|---|--|
| PE2(config)#evpn mpls enable | Activate the EVPN MPLS functionality on PE2, enabling it to participate in EVPN MPLS services. |
| PE2(config)#commit | Commit candidate configuration to be running configuration. |
| PE2(config)#evpn mpls vtep-ip-global 10.10.10.2 | Configure the global VTEP IP address 10.10.10.2, associating it with the loopback IP. |
| PE2(config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter for EVPN MPLS multi-homing. |

| | |
|--|--|
| PE2(config)#evpn mpls multihoming enable | Activate the EVPN MPLS multi-homing functionality, allowing PE2 to support multi-homed EVPN MPLS services. |
| PE2(config)#commit | Commit the transaction. |

PE2: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE2 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|---------------------------------------|--|
| PE2(config)#interface xe4 | Enter interface mode xe4. |
| PE2(config-if)#ip address 10.1.3.1/30 | Configure an IP address, 10.1.3.1/30, on the interface xe4. |
| PE2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE2(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe4 for area 1. |
| PE2(config-if)#exit | Exit interface mode xe4. |
| PE2(config)#commit | Commit the transaction. |
| PE2(config)#interface xe5 | Enter interface mode xe5. |
| PE2(config-if)#ip address 10.1.4.1/30 | Configure an IP address, 10.1.4.1/30, on the interface xe5. |
| PE2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE2(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe5 for area 1. |
| PE2(config-if)#exit | Exit interface mode xe5. |
| PE2(config)#commit | Commit the transaction. |

PE2: ISIS Configuration

The below configuration is performed to set up ISIS on PE2, to enable MPLS Traffic Engineering, Segment Routing, and other related features for efficient routing and network management.

| | |
|--------------------------------------|--|
| PE2(config)#router isis 1 | Enter router ISIS mode. |
| PE2(config-router)#is-type level-1-2 | Configure IS-Type as Level-1-2 specifies that the router will participate in both Level-1 and Level-2 areas within the ISIS network. |
| PE2(config-router)#metric-style wide | Configure the new style of metric type as wide. |

| | |
|--|--|
| PE2(config-router)#mpls traffic-eng router-id 10.10.10.2 | Configure the router's MPLS Traffic Engineering (TE) router ID TLV to 10.10.10.2, which is used for MPLS-TE path calculations. |
| PE2(config-router)#mpls traffic-eng level-1 | Enable MPLS-TE for IS-Type Level-1 routing. |
| PE2(config-router)#mpls traffic-eng level-2 | Enable MPLS-TE for IS-Type Level-2 routing. |
| PE2(config-router)#capability cspf | Enable Constraint Shortest Path First (CSPF) computation for traffic engineering. |
| PE2(config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| PE2(config-router)#fast-reroute ti-lfa level-1 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| PE2(config-router)#fast-reroute ti-lfa level-2 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| PE2(config-router)#bfd all-interfaces | Configure BFD on all interfaces for fast link failure detection. |
| PE2(config-router)#net 49.0000.0000.0002.00 | Set a Network Entity Title (NET) for this ISIS instance, specifying the area address and the system ID. |
| PE2(config-router)#isis segment-routing global block 17000 23500 | Enable ISIS SR globally and allocates label blocks for Segment Routing. |
| PE2(config-router)#segment-routing mpls | Enable SR ISIS. |
| PE2(config-router)#exit | Exit router ISIS mode and return to configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: BGP Configuration

The below BGP configuration on PE2 is established to enable BGP routing with ASN 65010, set the BGP router ID, define iBGP neighbors, configure BFD, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|---|--|
| PE2(config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| PE2(config-router)#bgp router-id 10.10.10.2 | Configure BGP router ID 10.10.10.2, identifying PE2 within the BGP network. |
| PE2(config-router)#neighbor 10.10.10.1 remote-as 65010 | Configure neighbor 10.10.10.1 as an iBGP neighbor with their remote AS number 65010. |
| PE2(config-router)#neighbor 10.10.10.1 update-source lo | Configure neighbor 10.10.10.1 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |

| | |
|---|---|
| PE2(config-router)#neighbor 10.10.10.3 remote-as 65010 | Configure neighbor 10.10.10.3 as an iBGP neighbor with their remote AS number 65010. |
| PE2(config-router)#neighbor 10.10.10.3 update-source lo | Configure neighbor 10.10.10.3 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE2(config-router)#neighbor 10.10.10.4 remote-as 65010 | Configure neighbor 10.10.10.4 as an iBGP neighbor with their remote AS number 65010. |
| PE2(config-router)#neighbor 10.10.10.4 update-source lo | Configure neighbor 10.10.10.4 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE2(config-router)#neighbor 10.10.10.1 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE2(config-router)#neighbor 10.10.10.3 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE2(config-router)#neighbor 10.10.10.4 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE2(config-router)#neighbor 10.10.10.1 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE2(config-router)#neighbor 10.10.10.3 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE2(config-router)#neighbor 10.10.10.4 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE2(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE2(config-router-af)#neighbor 10.10.10.1 activate | Activate EVPN for iBGP neighbor 10.10.10.1 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE2(config-router-af)#neighbor 10.10.10.3 activate | Activate EVPN for iBGP neighbor 10.10.10.3 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE2(config-router-af)#neighbor 10.10.10.4 activate | Activate EVPN for iBGP neighbor 10.10.10.4 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE2(config-router-af)#exit | Exit address family mode and return to the router BGP mode. |
| PE2(config-router)#commit | Commit the transaction. |
| PE2(config-router)#exit | Exit router BGP mode and return to the configure mode. |

PE2: MAC VRF Configuration

The below MAC VRF configuration on PE2 is carried out to define and set up VRFs named `vrf2` and `vp1s1001` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|---|--|
| PE2(config)#mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE2(config-vrf)#rd 10.10.10.2:1700 | Configure Route-Distinguisher value of 10.10.10.2:1700. |
| PE2(config-vrf)#route-target both 1700:1700 | Configure import and export values for the <code>vrf2</code> as 1700:1700. |
| PE2(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE2(config)#mac vrf vpls1001 | Enter VRF mode named <code>vpls1001</code> . |
| PE2(config-vrf)#rd 10.10.10.2:1001 | Configure Route-Distinguisher value of 10.10.10.2:1001. |
| PE2(config-vrf)#route-target both 1001:1001 | Configure import and export values for the <code>vpls1001</code> as 1001:1001. |
| PE2(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: EVPN and VRF Mapping

The below EVPN and VRF mapping configuration on PE2 is performed to establish mappings between EVPN identifiers and VRFs, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|--|--|
| PE2(config)#evpn mpls id 1800 xconnect target-mpls-id 1700 | Configure the EVPN-VPWS identifier with a source identifier of 1800 and a target identifier of 1700. |
| PE2(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Map VRF <code>vrf2</code> to the EVPN-VPWS identifier |
| PE2(config-evpn-mpls)#commit | Commit the transaction. |
| PE2(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| PE2(config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier an identifier of 3000. |
| PE2(config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Map VRF <code>vpls1001</code> to the EVPN-VPWS identifier |
| PE2(config-evpn-mpls)#commit | Commit the transaction. |
| PE2(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: Access Port Configuration for Port-active

The below access port configuration for port-active mode on PE2 is carried out to configure various parameters including system-mac, load balancing, service carving preferences, and EVPN settings for efficient network access and connectivity.

| | |
|---------------------------------|--|
| PE2(config)#interface po1 | Enter the port channel interface mode for <code>po1</code> |
| PE2(config-if)#load-interval 30 | Set the load interval to 30. |

| | |
|--|--|
| PE2(config-if)#evpn multi-homed system-mac 0000.1111.7777 load-balancing port-active | Configure the system-mac address 0000.1111.7777 for port-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE2(config-if-es)#service-carving auto | Configure service carving as auto, allowing automatic determination of service distribution preferences. |
| PE2(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE2(config-if)#exit | Exit interface mode po1 and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |
| PE2(config)#interface po1.1 switchport | Create a Layer 2 sub-interface po1.1 within the port channel. |
| PE2(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE2(config-if)#load-interval 30 | Set the load interval to 30. |
| PE2(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE2(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE2(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE2(config-if)#exit | Exit interface mode po1.1 and return to the configure mode. |
| PE2(config)#interface xe8 | Enter the interface mode for xe8. |
| PE2(config-if)#speed 10g | Set the speed to 10g. |
| PE2(config-if)#channel-group 1 mode active | Attach LAG interface po1. |
| PE2(config-if)#exit | Exit interface mode xe8 and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: Access Port Configuration for Single-active

The below access port configuration for single-active mode on PE2 is implemented to set up various parameters, including Ethernet Segment Identifier (ESI) settings, service carving preferences, and EVPN configurations, ensuring efficient network access and connectivity.

| | |
|--|--|
| PE2(config)#interface sa2 | Enter the single active interface mode for sa2. |
| PE2(config-if)#load-interval 30 | Set the load interval to 30. |
| PE2(config-if)#evpn multi-homed esi 00:00:11:11:77:77 load-balancing single-active | Configure the ESI with the value with the value 00:00:11:11:77:77 for single-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |

| | |
|---|--|
| PE2(config-if-es)#service-carving auto | Configure service carving as <code>auto</code> , allowing automatic determination of service distribution preferences. |
| PE2(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE2(config-if)#exit | Exit interface mode <code>sa2</code> and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |
| PE2(config)#interface sa2.1 switchport | Create a Layer 2 sub-interface <code>sa2.1</code> within the port channel. |
| PE2(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE2(config-if)#load-interval 30 | Set the load interval to 30. |
| PE2(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE2(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE2(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE2(config-if)#exit | Exit interface mode <code>sa2.1</code> and return to the configure mode. |
| PE2(config)#interface xe8 | Enter the interface mode for <code>xe8</code> . |
| PE2(config-if)#speed 10g | Set the speed to 10g. |
| PE2(config-if)#static-channel-group 2 | Attach the static-channel-group 2, the LAG interface <code>sa2</code> to <code>xe8</code> . |
| PE2(config-if)#exit | Exit interface mode <code>xe8</code> and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE3

PE3: Loopback Interface

The configuration on PE3 for a loopback interface with IP address `10.10.10.3/32` secondary is set up to provide IP connectivity for the router.

| | |
|---|---|
| PE3#configure terminal | Enter configure mode. |
| PE3(config)#interface lo | Enter the interface mode for the loopback interface <code>lo</code> . |
| PE3(config-if)#ip address 10.10.10.3/32 secondary | Configure a secondary IP address, <code>10.10.10.3/32</code> , on the loopback interface. |
| PE3(config-if)#ip router isis 1 | Enable ISIS routing on a loopback interface <code>lo</code> for area 1. |

| | |
|-------------------------------------|--|
| PE3(config-if)#prefix-sid index 800 | Configure a prefix segment identifier (prefix-SID) index value as 800. |
| PE3(config-if)#exit | Exit interface mode lo. |
| PE3(config)#commit | Commit the transaction. |

PE3: Configure SR

The following configurations aim to activate Segment Routing (SR) on PE3 and make MPLS the preferred method for segment routing, optimizing routing efficiency.

| | |
|-------------------------------|---|
| PE3(config)#segment-routing | Configure segment routing on PE3 device. |
| PE3(config-sr)#mpls sr-prefer | Set MPLS as the preferred segment routing protocol over others. |
| PE3(config-sr)#exit | Exit the router SR mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: Global LDP

The configuration on PE3 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|--|---|
| PE3(config)#router ldp | Enter the Router LDP mode. |
| PE3(config-router)#router-id 10.10.10.3 | Set the router ID for LDP to 10.10.10.3. |
| PE2(config-router)#transport-address ipv4 10.10.10.3 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| PE3(config-router)#targeted-peer ipv4 10.10.10.1 | Configure targeted peer for LDP using IPv4 addresses. |
| PE3(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE3(config-router)#targeted-peer ipv4 10.10.10.2 | Configure targeted peer for LDP using IPv4 addresses. |
| PE3(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE3(config-router)#targeted-peer ipv4 10.10.10.4 | Configure targeted peer for LDP using IPv4 addresses. |
| PE3(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE3(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: Global EVPN MPLS Command

The configuration on PE3 for the Global EVPN MPLS, includes activating EVPN MPLS, defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multi-homing, and activating EVPN MPLS multi-homing functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|--|--|
| <code>PE3(config)#evpn mpls enable</code> | Activate the EVPN MPLS functionality on PE3, enabling it to participate in EVPN MPLS services. |
| <code>PE3(config)#commit</code> | Commit candidate configuration to be running configuration. |
| <code>PE3(config)#evpn mpls vtep-ip-global 10.10.10.3</code> | Configure the global VTEP IP address 10.10.10.3, associating it with the loopback IP. |
| <code>PE3(config)#hardware-profile filter evpn-mpls-mh enable</code> | Enable hardware-profile filter for EVPN MPLS multi-homing. |
| <code>PE3(config)#evpn mpls multihoming enable</code> | Activate the EVPN MPLS multi-homing functionality, allowing PE3 to support multi-homed EVPN MPLS services. |
| <code>PE3(config)#commit</code> | Commit the transaction. |

PE3: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE3 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|--|--|
| <code>PE3(config)#interface xe1</code> | Enter interface mode xe1. |
| <code>PE3(config-if)#ip address 10.1.6.1/30</code> | Configure an IP address, 10.1.6.1/30, on the interface xe1. |
| <code>PE3(config-if)#enable-ldp ipv4</code> | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| <code>PE3(config-if)#label-switching</code> | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| <code>PE3(config-if)#ip router isis 1</code> | Enable ISIS routing on an interface xe1 for area 1. |
| <code>PE3(config-if)#exit</code> | Exit interface mode xe1. |
| <code>PE3(config)#commit</code> | Commit the transaction. |
| <code>PE3(config)#interface xe5</code> | Enter interface mode xe5. |
| <code>PE3(config-if)#ip address 10.1.7.1/30</code> | Configure an IP address, 10.1.7.1/30, on the interface xe5. |
| <code>PE3(config-if)#enable-ldp ipv4</code> | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| <code>PE3(config-if)#label-switching</code> | Enable label switching on the interface to enable MPLS-based packet forwarding. |

| | |
|---------------------------------|---|
| PE3(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe5 for area 1. |
| PE3(config-if)#exit | Exit interface mode xe5. |
| PE3(config)#commit | Commit the transaction. |

PE3: ISIS Configuration

The below configuration is performed to set up ISIS on PE3, to enable MPLS Traffic Engineering, Segment Routing, and other related features for efficient routing and network management.

| | |
|--|--|
| PE3(config)#router isis 1 | Enter router ISIS mode. |
| PE3(config-router)#is-type level-1-2 | Configure IS-Type as Level-1-2 specifies that the router will participate in both Level-1 and Level-2 areas within the ISIS network. |
| PE3(config-router)#metric-style wide | Configure the new style of metric type as wide. |
| PE3(config-router)#mpls traffic-eng router-id 10.10.10.3 | Configure the router's MPLS Traffic Engineering (TE) router ID TLV to 10.10.10.3, which is used for MPLS-TE path calculations. |
| PE3(config-router)#mpls traffic-eng level-1 | Enable MPLS-TE for IS-Type Level-1 routing. |
| PE3(config-router)#mpls traffic-eng level-2 | Enable MPLS-TE for IS-Type Level-2 routing. |
| PE3(config-router)#capability cspf | Enable Constraint Shortest Path First (CSPF) computation for traffic engineering. |
| PE3(config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| PE3(config-router)#fast-reroute ti-lfa level-1 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| PE3(config-router)#fast-reroute ti-lfa level-2 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| PE3(config-router)#bfd all-interfaces | Configure BFD on all interfaces for fast link failure detection. |
| PE3(config-router)#net 49.0000.0000.0003.00 | Set a Network Entity Title (NET) for this ISIS instance, specifying the area address and the system ID. |
| PE3(config-router)#isis segment-routing global block 17000 23500 | Enable ISIS SR globally and allocates label blocks for Segment Routing. |
| PE3(config-router)#segment-routing mpls | Enable SR ISIS. |
| PE3(config-router)#exit | Exit router ISIS mode and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: BGP Configuration

The below BGP configuration on PE3 is established to enable BGP routing with ASN 65010, set the BGP router ID, define iBGP neighbors, configure BFD, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|---|---|
| PE3(config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| PE3(config-router)#bgp router-id 10.10.10.3 | Configure BGP router ID 10.10.10.3, identifying PE3 within the BGP network. |
| PE3(config-router)#neighbor 10.10.10.1 remote-as 65010 | Configure neighbor 10.10.10.1 as an iBGP neighbor with their remote AS number 65010. |
| PE3(config-router)#neighbor 10.10.10.1 update-source lo | Configure neighbor 10.10.10.1 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE3(config-router)#neighbor 10.10.10.2 remote-as 65010 | Configure neighbor 10.10.10.2 as an iBGP neighbor with their remote AS number 65010. |
| PE3(config-router)#neighbor 10.10.10.2 update-source lo | Configure neighbor 10.10.10.2 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE3(config-router)#neighbor 10.10.10.4 remote-as 65010 | Configure neighbor 10.10.10.4 as an iBGP neighbor with their remote AS number 65010. |
| PE3(config-router)#neighbor 10.10.10.4 update-source lo | Configure neighbor 10.10.10.4 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE3(config-router)#neighbor 10.10.10.1 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE3(config-router)#neighbor 10.10.10.2 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE3(config-router)#neighbor 10.10.10.4 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE3(config-router)#neighbor 10.10.10.1 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE3(config-router)#neighbor 10.10.10.2 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE3(config-router)#neighbor 10.10.10.4 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE3(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE3(config-router-af)#neighbor 10.10.10.1 activate | Activate EVPN for iBGP neighbor 10.10.10.1 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE3(config-router-af)#neighbor 10.10.10.2 activate | Activate EVPN for iBGP neighbor 10.10.10.2 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |

| | |
|--|---|
| PE3(config-router-af)#neighbor 10.10.10.4 activate | Activate EVPN for iBGP neighbor 10.10.10.4 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE3(config-router-af)#exit | Exit address family mode and return to the router BGP mode. |
| PE3(config-router)#commit | Commit the transaction. |
| PE3(config-router)#exit | Exit router BGP mode and return to the configure mode. |

PE3: MAC VRF Configuration

The below MAC VRF configuration on PE3 is carried out to define and set up VRFs named `vrf2` and `vpls1001` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|---|--|
| PE3(config)#mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE3(config-vrf)#rd 10.10.10.3:1700 | Configure Route-Distinguisher value of 10.10.10.3:1700. |
| PE3(config-vrf)#route-target both 1700:1700 | Configure import and export values for the <code>vrf2</code> as 1700:1700. |
| PE3(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE3(config)#mac vrf vpls1001 | Enter VRF mode named <code>vpls1001</code> . |
| PE3(config-vrf)#rd 10.10.10.3:1001 | Configure Route-Distinguisher value of 10.10.10.3:1001. |
| PE3(config-vrf)#route-target both 1001:1001 | Configure import and export values for the <code>vpls1001</code> as 1001:1001. |
| PE3(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: EVPN and VRF Mapping

The below EVPN and VRF mapping configuration on PE3 is performed to establish mappings between EVPN identifiers and VRFs, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|--|--|
| PE3(config)#evpn mpls id 1700 xconnect target-mpls-id 1800 | Configure the EVPN-VPWS identifier with a source identifier of 1700 and a target identifier of 1800. |
| PE3(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Map VRF <code>vrf2</code> to the EVPN-VPWS identifier |
| PE3(config-evpn-mpls)#commit | Commit the transaction. |
| PE3(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| PE3(config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier an identifier of 3000. |

| | |
|--|---|
| PE3(config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Map VRF vpls1001 to the EVPN-VPWS identifier |
| PE3(config-evpn-mpls)#commit | Commit the transaction. |
| PE3(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE3: Access Port Configuration for Port-active

The below access port configuration for port-active mode on PE3 is carried out to configure various parameters including system-MAC, load balancing, service carving preferences, and EVPN settings for efficient network access and connectivity.

| | |
|--|--|
| PE3(config)#interface po1 | Enter the port channel interface mode for po1 |
| PE3(config-if)#load-interval 30 | Set the load interval to 30. |
| PE3(config-if)#evpn multi-homed system-mac 0000.2222.7777 load-balancing port-active | Configure the system-mac address 0000.2222.7777 for port-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE3(config-if-es)#service-carving auto | Configure service carving as auto, allowing automatic determination of service distribution preferences. |
| PE3(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE3(config-if)#exit | Exit interface mode po1 and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |
| PE3(config)#interface po1.1 switchport | Create a Layer 2 sub-interface po1.1 within the port channel. |
| PE3(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE3(config-if)#load-interval 30 | Set the load interval to 30. |
| PE3(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE3(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE3(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE3(config-if)#exit | Exit interface mode po1.1 and return to the configure mode. |
| PE3(config)#interface xe2 | Enter the interface mode for xe2. |
| PE3(config-if)#speed 10g | Set the speed to 10g. |
| PE3(config-if)#channel-group 1 mode active | Attach LAG interface po1. |
| PE3(config-if)#exit | Exit interface mode xe2 and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: Access Port Configuration for Single-active

The below access port configuration for single-active mode on PE3 is implemented to set up various parameters, including Ethernet Segment Identifier (ESI) settings, service carving preferences, and EVPN configurations, ensuring efficient network access and connectivity.

| | |
|---|--|
| PE3(config)#interface sa1 | Enter the single active interface mode for sa1 . |
| PE3(config-if)#load-interval 30 | Set the load interval to 30. |
| PE3(config-if)#evpn multi-homed esi 00:00:22:22:77:77 load-balancing single-active | Configure the ESI with the value with the value 00:00:22:22:77:77 for single-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE3(config-if-es)#service-carving auto | Configure service carving as auto, allowing automatic determination of service distribution preferences. |
| PE3(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE3(config-if)#exit | Exit interface mode sa1 and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |
| PE3(config)#interface sa1.1 switchport | Create a Layer 2 sub-interface sa1.1 within the port channel. |
| PE3(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE3(config-if)#load-interval 30 | Set the load interval to 30. |
| PE3(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE3(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE3(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE3(config-if)#exit | Exit interface mode sa1.1 and return to the configure mode. |
| PE3(config)#interface xe2 | Enter the interface mode for xe2 . |
| PE3(config-if)#speed 10g | Set the speed to 10g. |
| PE3(config-if)#static-channel-group 1 | Attach the static-channel-group 1, the LAG interface sa1 to xe2. |
| PE3(config-if)#exit | Exit interface mode xe2 and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE4

PE4: Loopback Interface

The configuration on PE4 for a loopback interface with IP address 10.10.10.4/32 secondary is set up to provide IP connectivity for the router.

| | |
|---|---|
| PE4#configure terminal | Enter configure mode. |
| PE4(config)#interface lo | Enter the interface mode for the loopback interface lo. |
| PE4(config-if)#ip address 10.10.10.4/32 secondary | Configure a secondary IP address, 10.10.10.4/32, on the loopback interface. |
| PE4(config-if)#ip router isis 1 | Enable ISIS routing on a loopback interface lo for area 1. |
| PE4(config-if)#prefix-sid index 800 | Configure a prefix segment identifier (prefix-SID) index value as 800. |
| PE4(config-if)#exit | Exit interface mode lo. |
| PE4(config)#commit | Commit the transaction. |

PE4: Configure SR

The following configurations aim to activate Segment Routing (SR) on PE4 and make MPLS the preferred method for segment routing, optimizing routing efficiency.

| | |
|-------------------------------|---|
| PE4(config)#segment-routing | Configure segment routing on PE4 device. |
| PE4(config-sr)#mpls sr-prefer | Set MPLS as the preferred segment routing protocol over others. |
| PE4(config-sr)#exit | Exit the router SR mode. |
| PE4(config)#commit | Commit the transaction. |

PE4: Global LDP

The configuration on PE4 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|--|---|
| PE4(config)#router ldp | Enter the Router LDP mode. |
| PE4(config-router)#router-id 10.10.10.4 | Set the router ID for LDP to 10.10.10.4. |
| PE4(config-router)#transport-address ipv4 10.10.10.4 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| PE4(config-router)#targeted-peer ipv4 10.10.10.1 | Configure targeted peer for LDP using IPv4 addresses. |

| | |
|---|--|
| PE4(config-router-targeted-peer) #exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE4(config-router) #targeted-peer ipv4 10.10.10.2 | Configure targeted peer for LDP using IPv4 addresses. |
| PE4(config-router-targeted-peer) #exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE4(config-router) #targeted-peer ipv4 10.10.10.3 | Configure targeted peer for LDP using IPv4 addresses. |
| PE4(config-router-targeted-peer) #exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE4(config-router) #exit | Exit router LDP mode and return to the configure mode. |
| PE4(config) #commit | Commit the transaction. |

PE4: Global EVPN MPLS Command

The configuration on PE4 for the Global EVPN MPLS, includes activating EVPN MPLS, defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multi-homing, and activating EVPN MPLS multi-homing functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|--|--|
| PE4(config) #evpn mpls enable | Activate the EVPN MPLS functionality on PE4, enabling it to participate in EVPN MPLS services. |
| PE4(config) #commit | Commit candidate configuration to be running configuration. |
| PE4(config) #evpn mpls vtep-ip-global 10.10.10.4 | Configure the global VTEP IP address 10.10.10.4, associating it with the loopback IP. |
| PE4(config) #hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter for EVPN MPLS multi-homing. |
| PE4(config) #evpn mpls multihoming enable | Activate the EVPN MPLS multi-homing functionality, allowing PE4 to support multi-homed EVPN MPLS services. |
| PE4(config) #commit | Commit the transaction. |

PE4: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE4 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|--|--|
| PE4(config) #interface xe2 | Enter interface mode xe2. |
| PE4(config-if) #ip address 10.1.5.1/30 | Configure an IP address, 10.1.5.1/30, on the interface xe2. |
| PE4(config-if) #enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE4(config-if) #label-switching | Enable label switching on the interface to enable |

| | |
|---------------------------------------|--|
| | MPLS-based packet forwarding. |
| PE4(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe2 for area 1. |
| PE4(config-if)#exit | Exit interface mode xe2. |
| PE4(config)#commit | Commit the transaction. |
| PE4(config)#interface xe0 | Enter interface mode xe0. |
| PE4(config-if)#ip address 10.1.8.1/30 | Configure an IP address, 10.1.8.1/30, on the interface xe0. |
| PE4(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE4(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE4(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe2 for area 1. |
| PE4(config-if)#exit | Exit interface mode xe0. |
| PE4(config)#commit | Commit the transaction. |

PE4: ISIS Configuration

The below configuration is performed to set up ISIS on PE4, to enable MPLS Traffic Engineering, Segment Routing, and other related features for efficient routing and network management.

| | |
|---|--|
| PE4(config)#router isis 1 | Enter router ISIS mode. |
| PE4(config-router)#is-type level-1-2 | Configure IS-Type as Level-1-2 specifies that the router will participate in both Level-1 and Level-2 areas within the ISIS network. |
| PE4(config-router)#metric-style wide | Configure the new style of metric type as wide. |
| PE4(config-router)#mpls traffic-eng router-id 10.10.10.4 | Configure the router's MPLS Traffic Engineering (TE) router ID TLV to 10.10.10.4, which is used for MPLS-TE path calculations. |
| PE4(config-router)#mpls traffic-eng level-1 | Enable MPLS-TE for IS-Type Level-1 routing. |
| PE4(config-router)#mpls traffic-eng level-2 | Enable MPLS-TE for IS-Type Level-2 routing. |
| PE4(config-router)#capability cspf | Enable Constraint Shortest Path First (CSPF) computation for traffic engineering. |
| PE4(config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| PE4(config-router)#fast-reroute ti-lfa level-1 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| PE4(config-router)#fast-reroute ti-lfa level-2 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR |

| | |
|--|---|
| | inequalities |
| PE4(config-router)#bfd all-interfaces | Configure BFD on all interfaces for fast link failure detection. |
| PE4(config-router)#net 49.0000.0000.0004.00 | Set a Network Entity Title (NET) for this ISIS instance, specifying the area address and the system ID. |
| PE4(config-router)#isis segment-routing global block 17000 23500 | Enable ISIS SR globally and allocates label blocks for Segment Routing. |
| PE4(config-router)#segment-routing mpls | Enable SR ISIS. |
| PE4(config-router)#exit | Exit router ISIS mode and return to configure mode. |
| PE4(config)#commit | Commit the transaction. |

PE4: BGP Configuration

The below BGP configuration on PE4 is established to enable BGP routing with ASN 65010, set the BGP router ID, define iBGP neighbors, configure BFD, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|---|--|
| PE4(config)#router bgp 65010 | Enter the Router BGP mode, ASN: 65010 |
| PE4(config-router)#bgp router-id 10.10.10.4 | Configure BGP router ID 10.10.10.4, identifying PE4 within the BGP network. |
| PE4(config-router)#neighbor 10.10.10.1 remote-as 65010 | Configure neighbor 10.10.10.1 as an iBGP neighbor with their remote AS number 65010. |
| PE4(config-router)#neighbor 10.10.10.1 update-source lo | Configure neighbor 10.10.10.1 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE4(config-router)#neighbor 10.10.10.2 remote-as 65010 | Configure neighbor 10.10.10.2 as an iBGP neighbor with their remote AS number 65010. |
| PE4(config-router)#neighbor 10.10.10.2 update-source lo | Configure neighbor 10.10.10.2 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE4(config-router)#neighbor 10.10.10.3 remote-as 65010 | Configure neighbor 10.10.10.3 as an iBGP neighbor with their remote AS number 65010. |
| PE4(config-router)#neighbor 10.10.10.3 update-source lo | Configure neighbor 10.10.10.3 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE4(config-router)#neighbor 10.10.10.1 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE4(config-router)#neighbor 10.10.10.2 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |
| PE4(config-router)#neighbor 10.10.10.3 fall-over bfd multihop | Configure BFD for the BGP neighbor to provide rapid failure detection. |

| | |
|--|---|
| PE4(config-router)#neighbor 10.10.10.1 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE4(config-router)#neighbor 10.10.10.2 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE4(config-router)#neighbor 10.10.10.3 advertisement-interval 0 | Configure advertisement interval for the neighbor, allowing more frequent BGP updates. |
| PE4(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE4(config-router-af)#neighbor 10.10.10.1 activate | Activate EVPN for iBGP neighbor 10.10.10.1 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE4(config-router-af)#neighbor 10.10.10.2 activate | Activate EVPN for iBGP neighbor 10.10.10.2 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE4(config-router-af)#neighbor 10.10.10.3 activate | Activate EVPN for iBGP neighbor 10.10.10.3 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE4(config-router-af)#exit | Exit address family mode and return to the router BGP mode. |
| PE4(config-router)#commit | Commit the transaction. |
| PE4(config-router)#exit | Exit router BGP mode and return to the configure mode. |

PE4: MAC VRF Configuration

The below MAC VRF configuration on PE4 is carried out to define and set up VRFs named `vrf2` and `vpls1001` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|---|--|
| PE4(config)#mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE4(config-vrf)#rd 10.10.10.4:1700 | Configure Route-Distinguisher value of 10.10.10.4:1700. |
| PE4(config-vrf)#route-target both 1700:1700 | Configure import and export values for the <code>vrf2</code> as 1700:1700. |
| PE4(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE4(config)#mac vrf vpls1001 | Enter VRF mode named <code>vpls1001</code> . |
| PE4(config-vrf)#rd 10.10.10.4:1001 | Configure Route-Distinguisher value of 10.10.10.4:1001. |
| PE4(config-vrf)#route-target both 1001:1001 | Configure import and export values for the <code>vpls1001</code> as 1001:1001. |
| PE4(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |

PE4: EVPN and VRF Mapping

The below EVPN and VRF mapping configuration on PE4 is performed to establish mappings between EVPN identifiers and VRFs, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|--|--|
| PE4(config)#evpn mpls id 1700 xconnect target-mpls-id 1800 | Configure the EVPN-VPWS identifier with a source identifier of 1700 and a target identifier of 1800. |
| PE4(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Map VRF <code>vrf2</code> to the EVPN-VPWS identifier |
| PE4(config-evpn-mpls)#commit | Commit the transaction. |
| PE4(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| PE4(config)#evpn mpls id 3000 | Configure the EVPN-VPLS identifier an identifier of 3000. |
| PE4(config-evpn-mpls)#host-reachability-protocol evpn-bgp vpls1001 | Map VRF <code>vpls1001</code> to the EVPN-VPWS identifier |
| PE4(config-evpn-mpls)#commit | Commit the transaction. |
| PE4(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE4: Access Port Configuration for Port-active

The below access port configuration for port-active mode on PE4 is carried out to configure various parameters including system-MAC, load balancing, service carving preferences, and EVPN settings for efficient network access and connectivity.

| | |
|--|--|
| PE4(config)#interface po1 | Enter the port channel interface mode for <code>po1</code> |
| PE4(config-if)#load-interval 30 | Set the load interval to 30. |
| PE4(config-if)#evpn multi-homed system-mac 0000.2222.7777 load-balancing port-active | Configure the system-mac address 0000.2222.7777 for port-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE4(config-if-es)#service-carving auto | Configure service carving as <code>auto</code> , allowing automatic determination of service distribution preferences. |
| PE4(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE4(config-if)#exit | Exit interface mode <code>po1</code> and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |
| PE4(config)#interface po1.1 switchport | Create a Layer 2 sub-interface <code>po1.1</code> within the port channel. |
| PE4(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |

| | |
|--|--|
| PE4(config-if)#load-interval 30 | Set the load interval to 30. |
| PE4(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE4(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |
| PE4(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE4(config-if)#exit | Exit interface mode <code>po1.1</code> and return to the configure mode. |
| PE4(config)#interface xe11 | Enter the interface mode for <code>xe11</code> . |
| PE4(config-if)#speed 10g | Set the speed to 10g. |
| PE4(config-if)#channel-group 1 mode active | Attach LAG interface <code>po1</code> . |
| PE4(config-if)#exit | Exit interface mode <code>xe11</code> and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |

PE4: Access Port Configuration for Single-active

The below access port configuration for single-active mode on PE4 is implemented to set up various parameters, including Ethernet Segment Identifier (ESI) settings, service carving preferences, and EVPN configurations, ensuring efficient network access and connectivity.

| | |
|---|---|
| PE4(config)#interface sa2 | Enter the single active interface mode for <code>sa2</code> . |
| PE4(config-if)#load-interval 30 | Set the load interval to 30. |
| PE4(config-if)#evpn multi-homed esi 00:00:22:22:77:77 load-balancing single-active | Configure the ESI with the value with the value <code>00:00:22:22:77:77</code> for single-active mode, which plays a role in load balancing and enter to the EVPN Ethernet Segment (ES) mode. |
| PE4(config-if-es)#service-carving auto | Configure service carving as <code>auto</code> , allowing automatic determination of service distribution preferences. |
| PE4(config-if-es)#exit | Exit the EVPN ES mode and return to the configure mode. |
| PE4(config-if)#exit | Exit interface mode <code>sa2</code> and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |
| PE4(config)#interface sa2.1 switchport | Create a Layer 2 sub-interface <code>sa2.1</code> within the port channel. |
| PE4(config-if)#encapsulation dot1q 100 | Set encapsulation to dot1q with VLAN ID 100. |
| PE4(config-if)#load-interval 30 | Set the load interval to 30. |
| PE4(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE4(config-acc-if-evpn)#map vpn-id 1800 | Map VPN-ID 1800. |

| | |
|---------------------------------------|--|
| PE4(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE4(config-if)#exit | Exit interface mode <code>sa2.1</code> and return to the configure mode. |
| PE4(config)#interface xe11 | Enter the interface mode for <code>xe11</code> . |
| PE4(config-if)#speed 10g | Set the speed to 10g. |
| PE4(config-if)#static-channel-group 2 | Attach the static-channel-group 2, the LAG interface <code>sa2</code> to <code>xe11</code> . |
| PE4(config-if)#exit | Exit interface mode <code>xe11</code> and return to the configure mode. |
| PE4(config)#commit | Commit the transaction. |

Core Routers Configurations

P1: Loopback Interface

The configuration on P1 for a loopback interface with IP address `10.10.10.5/32` secondary is set up to provide IP connectivity for the router.

| | |
|--|---|
| P1#configure terminal | Enter configure mode. |
| P1(config)#interface lo | Enter the interface mode for the loopback interface <code>lo</code> . |
| P1(config-if)#ip address 10.10.10.5/32 secondary | Configure a secondary IP address, <code>10.10.10.5/32</code> , on the loopback interface. |
| P1(config-if)#ip router isis 1 | Enable ISIS routing on a loopback interface <code>lo</code> for area 1. |
| P1(config-if)#prefix-sid index 800 | Configure a prefix segment identifier (prefix-SID) index value as 800. |
| P1(config-if)#exit | Exit interface mode <code>lo</code> . |
| P1(config)#commit | Commit the transaction. |

P1: Configure SR

The following configurations aim to activate Segment Routing (SR) on P1 and make MPLS the preferred method for segment routing, optimizing routing efficiency.

| | |
|------------------------------|---|
| P1(config)#segment-routing | Configure segment routing on P1 device. |
| P1(config-sr)#mpls sr-prefer | Set MPLS as the preferred segment routing protocol over others. |
| P1(config-sr)#exit | Exit the router SR mode. |
| P1(config)#commit | Commit the transaction. |

P1: Global LDP

The configuration on P1 for the Global LDP router, specifying router ID and targeted peer, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|---|---|
| P1(config)#router ldp | Enter the Router LDP mode. |
| P1(config-router)#router-id 10.10.10.5 | Set the router ID for LDP to 10.10.10.5. |
| P1(config-router)#transport-address ipv4 10.10.10.5 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| P1(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| P1(config)#commit | Commit the transaction. |

P1: Interface Configuration

The below configuration is performed to set up interfaces on P1 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|--------------------------------------|--|
| P1(config)#interface xe1 | Enter interface mode xe1. |
| P1(config-if)#ip address 10.1.1.2/30 | Configure an IP address, 10.1.1.2/30, on the interface xe1. |
| P1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P1(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe1 for area 1. |
| P1(config-if)#exit | Exit interface mode xe1. |
| P1(config)#commit | Commit the transaction. |
| P1(config)#interface xe2 | Enter interface mode xe2. |
| P1(config-if)#ip address 10.1.4.2/30 | Configure an IP address, 10.1.4.2/30, on the interface xe2. |
| P1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P1(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe2 for area 1. |
| P1(config-if)#exit | Exit interface mode xe2. |
| P1(config)#commit | Commit the transaction. |

| | |
|--------------------------------------|--|
| P1(config)#interface xe3 | Enter interface mode xe3. |
| P1(config-if)#ip address 10.1.5.2/30 | Configure an IP address, 10.1.5.2/30, on the interface xe3. |
| P1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P1(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe3 for area 1. |
| P1(config-if)#exit | Exit interface mode xe3. |
| P1(config)#commit | Commit the transaction. |
| P1(config)#interface xe4 | Enter interface mode xe4. |
| P1(config-if)#ip address 10.1.6.2/30 | Configure an IP address, 10.1.6.2/30, on the interface xe4. |
| P1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P1(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe4 for area 1. |
| P1(config-if)#exit | Exit interface mode xe4. |
| P1(config)#commit | Commit the transaction. |

P1: ISIS Configuration

The below configuration is performed to set up ISIS on P1, to enable MPLS Traffic Engineering, Segment Routing, and other related features for efficient routing and network management.

| | |
|---|--|
| P1(config)#router isis 1 | Enter router ISIS mode. |
| P1(config-router)#is-type level-1-2 | Configure IS-Type as Level-1-2 specifies that the router will participate in both Level-1 and Level-2 areas within the ISIS network. |
| P1(config-router)#metric-style wide | Configure the new style of metric type as wide. |
| P1(config-router)#mpls traffic-eng router-id 10.10.10.5 | Configure the router's MPLS Traffic Engineering (TE) router ID TLV to 10.10.10.5, which is used for MPLS-TE path calculations. |
| P1(config-router)#mpls traffic-eng level-1 | Enable MPLS-TE for IS-Type Level-1 routing. |
| P1(config-router)#mpls traffic-eng level-2 | Enable MPLS-TE for IS-Type Level-2 routing. |
| P1(config-router)#capability cspf | Enable Constraint Shortest Path First (CSPF) computation for traffic engineering. |
| P1(config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |

| | |
|---|--|
| P1(config-router)#fast-reroute ti-lfa level-1 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| P1(config-router)#fast-reroute ti-lfa level-2 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| P1(config-router)#bfd all-interfaces | Configure BFD on all interfaces for fast link failure detection. |
| P1(config-router)#net 49.0000.0000.0005.00 | Set a Network Entity Title (NET) for this ISIS instance, specifying the area address and the system ID. |
| P1(config-router)#isis segment-routing global block 17000 23500 | Enable ISIS SR globally and allocates label blocks for Segment Routing. |
| P1(config-router)#segment-routing mpls | Enable SR ISIS. |
| P1(config-router)#exit | Exit router ISIS mode and return to the configure mode. |
| P1(config)#commit | Commit the transaction. |

P2: Loopback Interface

The configuration on P2 for a loopback interface with IP address 10.10.10.6/32 secondary is set up to provide IP connectivity for the router.

| | |
|--|---|
| P2#configure terminal | Enter configure mode. |
| P2(config)#interface lo | Enter the interface mode for the loopback interface lo. |
| P2(config-if)#ip address 10.10.10.6/32 secondary | Configure a secondary IP address, 10.10.10.6/32, on the loopback interface. |
| P2(config-if)#ip router isis 1 | Enable ISIS routing on a loopback interface lo for area 1. |
| P2(config-if)#prefix-sid index 800 | Configure a prefix segment identifier (prefix-SID) index value as 800. |
| P2(config-if)#exit | Exit interface mode lo. |
| P2(config)#commit | Commit the transaction. |

P2: Configure SR

The following configurations aim to activate Segment Routing (SR) on P2 and make MPLS the preferred method for segment routing, optimizing routing efficiency.

| | |
|------------------------------|---|
| P2(config)#segment-routing | Configure segment routing on P2 device. |
| P2(config-sr)#mpls sr-prefer | Set MPLS as the preferred segment routing protocol over others. |

| | |
|--------------------|--------------------------|
| P2(config-sr)#exit | Exit the router SR mode. |
| P2(config)#commit | Commit the transaction. |

P2: Global LDP

The configuration on P2 for the Global LDP router, specifying router ID and targeted peer, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|---|---|
| P2(config)#router ldp | Enter the Router LDP mode. |
| P2(config-router)#router-id 10.10.10.6 | Set the router ID for LDP to 10.10.10.6. |
| P2(config-router)#transport-address ipv4 10.10.10.6 | Configure the transport address for IPv4 (for IPv6 use ipv6 parameter) to be used for a TCP session where LDP operates. Note: It is preferable to use the loopback address as the transport address. |
| P2(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| P2(config)#commit | Commit the transaction. |

P2: Interface Configuration

The below configuration is performed to set up interfaces on P2 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|--------------------------------------|--|
| P2(config)#interface xe12 | Enter interface mode xe12. |
| P2(config-if)#ip address 10.1.2.2/30 | Configure an IP address, 10.1.2.2/30, on the interface xe12. |
| P2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P2(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe12 for area 1. |
| P2(config-if)#exit | Exit interface mode xe12. |
| P2(config)#commit | Commit the transaction. |
| P2(config)#interface xe13 | Enter interface mode xe13. |
| P2(config-if)#ip address 10.1.3.2/30 | Configure an IP address, 10.1.3.2/30, on the interface xe13. |
| P2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P2(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe13 for area 1. |

| | |
|--------------------------------------|--|
| P2(config-if)#exit | Exit interface mode xe13. |
| P2(config)#commit | Commit the transaction. |
| P2(config)#interface xe11 | Enter interface mode xe11. |
| P2(config-if)#ip address 10.1.7.2/30 | Configure an IP address, 10.1.7.2/30, on the interface xe11. |
| P2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P2(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe11 for area 1. |
| P2(config-if)#exit | Exit interface mode xe11. |
| P2(config)#commit | Commit the transaction. |
| P2(config)#interface xe14 | Enter interface mode xe14. |
| P2(config-if)#ip address 10.1.8.2/30 | Configure an IP address, 10.1.8.2/30, on the interface xe14. |
| P2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P2(config-if)#ip router isis 1 | Enable ISIS routing on an interface xe14 for area 1. |
| P2(config-if)#exit | Exit interface mode xe14. |
| P2(config)#commit | Commit the transaction. |

P2: ISIS Configuration

The below configuration is performed to set up ISIS on P2, to enable MPLS Traffic Engineering, Segment Routing, and other related features for efficient routing and network management.

| | |
|---|--|
| P2(config)#router isis 1 | Enter router ISIS mode. |
| P2(config-router)#is-type level-1-2 | Configure IS-Type as Level-1-2 specifies that the router will participate in both Level-1 and Level-2 areas within the ISIS network. |
| P2(config-router)#metric-style wide | Configure the new style of metric type as wide. |
| P2(config-router)#mpls traffic-eng router-id 10.10.10.6 | Configure the router's MPLS Traffic Engineering (TE) router ID TLV to 10.10.10.6, which is used for MPLS-TE path calculations. |
| P2(config-router)#mpls traffic-eng level-1 | Enable MPLS-TE for IS-Type Level-1 routing. |
| P2(config-router)#mpls traffic-eng level-2 | Enable MPLS-TE for IS-Type Level-2 routing. |
| P2(config-router)#capability cspf | Enable Constraint Shortest Path First (CSPF) computation for traffic engineering. |

| | |
|---|--|
| P2(config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| P2(config-router)#fast-reroute ti-lfa level-1 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| P2(config-router)#fast-reroute ti-lfa level-2 proto ipv4 | Configure Remote Loop-Free Alternate (LFA) to calculate backup paths to those destinations whichever does not satisfy basic LFA FRR inequalities |
| P2(config-router)#bfd all-interfaces | Configure BFD on all interfaces for fast link failure detection. |
| P2(config-router)#net 49.0000.0000.0006.00 | Set a Network Entity Title (NET) for this ISIS instance, specifying the area address and the system ID. |
| P2(config-router)#isis segment-routing global block 17000 23500 | Enable ISIS SR globally and allocates label blocks for Segment Routing. |
| P2(config-router)#segment-routing mpls | Enable SR ISIS. |
| P2(config-router)#exit | Exit router ISIS mode and return to the configure mode. |
| P2(config)#commit | Commit the transaction. |

EVPN SR Active-Standby MH Validation

This section provides show outputs validation for port-active mode, covering ELINE and ELAN services with SR as the underlay MPLS path.

The following show output displays the forwarding table entries on PE1, PE2, PE3, and PE4 devices in the network [Figure 55](#) using the **show mpls forwarding-table** command.

```
PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP
```

```
FTN-ECMP LDP: Disabled
Code  FEC          FTN-ID  Nhlfe-ID  Tunnel-id  Pri  LSP-Type  Out-Label  Out-
Intf  ELC    Nexthop
i>    10.10.10.2/32  1      4         0          Yes  LSP_
DEFAULT 17700      xe2      No        10.1.2.2
i>    10.10.10.2/32  10     23        0          No   LSP_
DEFAULT 3          xe1      No        10.1.1.2
i(b)>  10.10.10.2/32  6      17        2201      Yes  LSP_
DEFAULT 17700      xe1      No        10.1.1.2
i>    10.10.10.3/32  2      6         0          Yes  LSP_
DEFAULT 17400      xe2      No        10.1.2.2
i>    10.10.10.3/32  11     24        0          No   LSP_
DEFAULT 3          xe1      No        10.1.1.2
i(b)>  10.10.10.3/32  7      19        2202      Yes  LSP_
DEFAULT 17400      xe1      No        10.1.1.2
i(b)>  10.10.10.3/32  9      22        2204      Yes  LSP_
DEFAULT 17400      xe2      No        10.1.2.2
i>    10.10.10.4/32  3      8         0          Yes  LSP_
DEFAULT 17300      xe2      No        10.1.2.2
```



```

i> 10.10.10.4/32 12 25 0 No LSP_
DEFAULT 3 xe1 No 10.1.1.2
i(b)> 10.10.10.4/32 8 21 2203 Yes LSP_
DEFAULT 17300 xe1 No 10.1.1.2
i> 10.10.10.5/32 4 9 0 Yes LSP_
DEFAULT 3 xe1 No 10.1.1.2
i> 10.10.10.5/32 13 27 0 No LSP_
DEFAULT 17600 xe2 No 10.1.2.2
i> 10.10.10.6/32 5 15 0 Yes LSP_
DEFAULT 3 xe2 No 10.1.2.2
i> 10.10.10.6/32 14 29 0 No LSP_
DEFAULT 17500 xe1 No 10.1.1.2

```

PE2#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled

| Code | FEC | Nexthop | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|---------|--------|----------|-----------|-----|----------|-----------|------|
| Intf | ELC | | | | | | | | |
| i> | 10.10.10.1/32 | | 1 | 10 | 0 | Yes | LSP_ | | |
| DEFAULT | 17800 | xe4 | No | 10.1.3.2 | | | | | |
| i> | 10.10.10.1/32 | | 10 | 27 | 0 | No | LSP_ | | |
| DEFAULT | 3 | xe5 | No | 10.1.4.2 | | | | | |
| i(b)> | 10.10.10.1/32 | | 6 | 21 | 2201 | Yes | LSP_ | | |
| DEFAULT | 17800 | xe5 | No | 10.1.4.2 | | | | | |
| i> | 10.10.10.3/32 | | 2 | 11 | 0 | Yes | LSP_ | | |
| DEFAULT | 17400 | xe4 | No | 10.1.3.2 | | | | | |
| i> | 10.10.10.3/32 | | 11 | 28 | 0 | No | LSP_ | | |
| DEFAULT | 3 | xe5 | No | 10.1.4.2 | | | | | |
| i(b)> | 10.10.10.3/32 | | 7 | 23 | 2202 | Yes | LSP_ | | |
| DEFAULT | 17400 | xe5 | No | 10.1.4.2 | | | | | |
| i(b)> | 10.10.10.3/32 | | 9 | 26 | 2204 | Yes | LSP_ | | |
| DEFAULT | 17400 | xe4 | No | 10.1.3.2 | | | | | |
| i> | 10.10.10.4/32 | | 3 | 12 | 0 | Yes | LSP_ | | |
| DEFAULT | 17300 | xe4 | No | 10.1.3.2 | | | | | |
| i> | 10.10.10.4/32 | | 12 | 29 | 0 | No | LSP_ | | |
| DEFAULT | 3 | xe5 | No | 10.1.4.2 | | | | | |
| i(b)> | 10.10.10.4/32 | | 8 | 25 | 2203 | Yes | LSP_ | | |
| DEFAULT | 17300 | xe5 | No | 10.1.4.2 | | | | | |
| i> | 10.10.10.5/32 | | 4 | 13 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe5 | No | 10.1.4.2 | | | | | |
| i> | 10.10.10.5/32 | | 13 | 31 | 0 | No | LSP_ | | |
| DEFAULT | 17600 | xe4 | No | 10.1.3.2 | | | | | |
| i> | 10.10.10.6/32 | | 5 | 19 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe4 | No | 10.1.3.2 | | | | | |
| i> | 10.10.10.6/32 | | 14 | 33 | 0 | No | LSP_ | | |
| DEFAULT | 17500 | xe5 | No | 10.1.4.2 | | | | | |

PE3#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled

| Code | FEC | Nexthop | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|---------|--------|----------|-----------|-----|----------|-----------|------|
| Intf | ELC | | | | | | | | |
| i> | 10.10.10.1/32 | | 1 | 4 | 0 | Yes | LSP_ | | |
| DEFAULT | 17800 | xe5 | No | 10.1.7.2 | | | | | |
| i> | 10.10.10.1/32 | | 10 | 23 | 0 | No | LSP_ | | |
| DEFAULT | 3 | xe1 | No | 10.1.6.2 | | | | | |
| i(b)> | 10.10.10.1/32 | | 6 | 17 | 2201 | Yes | LSP_ | | |
| DEFAULT | 17800 | xe1 | No | 10.1.6.2 | | | | | |
| i(b)> | 10.10.10.1/32 | | 9 | 22 | 2204 | Yes | LSP_ | | |
| DEFAULT | 17800 | xe5 | No | 10.1.7.2 | | | | | |
| i> | 10.10.10.2/32 | | 2 | 6 | 0 | Yes | LSP_ | | |

```

DEFAULT 17700 xe5 No 10.1.7.2
i> 10.10.10.2/32 11 24 0 No LSP_
DEFAULT 3 xe1 No 10.1.6.2
i(b)> 10.10.10.2/32 7 19 2202 Yes LSP_
DEFAULT 17700 xe1 No 10.1.6.2
i> 10.10.10.4/32 3 8 0 Yes LSP_
DEFAULT 17300 xe5 No 10.1.7.2
i> 10.10.10.4/32 12 25 0 No LSP_
DEFAULT 3 xe1 No 10.1.6.2
i(b)> 10.10.10.4/32 8 21 2203 Yes LSP_
DEFAULT 17300 xe1 No 10.1.6.2
i> 10.10.10.5/32 4 9 0 Yes LSP_
DEFAULT 3 xe1 No 10.1.6.2
i> 10.10.10.5/32 13 27 0 No LSP_
DEFAULT 17600 xe5 No 10.1.7.2
i> 10.10.10.6/32 5 15 0 Yes LSP_
DEFAULT 3 xe5 No 10.1.7.2
i> 10.10.10.6/32 14 29 0 No LSP_
DEFAULT 17500 xe1 No 10.1.6.2

```

PE4#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
 B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled

| Code | FEC | Nexthop | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|---------|--------|----------|-----------|-----|----------|-----------|------|
| Intf | ELC | | | | | | | | |
| i> | 10.10.10.1/32 | | 1 | 4 | 0 | Yes | LSP_ | | |
| DEFAULT | 17800 | xe0 | No | 10.1.8.2 | | | | | |
| i> | 10.10.10.1/32 | | 10 | 23 | 0 | No | LSP_ | | |
| DEFAULT | 3 | xe2 | No | 10.1.5.2 | | | | | |
| i(b)> | 10.10.10.1/32 | | 6 | 17 | 2201 | Yes | LSP_ | | |
| DEFAULT | 17800 | xe2 | No | 10.1.5.2 | | | | | |
| i> | 10.10.10.2/32 | | 2 | 6 | 0 | Yes | LSP_ | | |
| DEFAULT | 17700 | xe0 | No | 10.1.8.2 | | | | | |
| i> | 10.10.10.2/32 | | 11 | 24 | 0 | No | LSP_ | | |
| DEFAULT | 3 | xe2 | No | 10.1.5.2 | | | | | |
| i(b)> | 10.10.10.2/32 | | 7 | 19 | 2202 | Yes | LSP_ | | |
| DEFAULT | 17700 | xe2 | No | 10.1.5.2 | | | | | |
| i> | 10.10.10.3/32 | | 3 | 8 | 0 | Yes | LSP_ | | |
| DEFAULT | 17400 | xe0 | No | 10.1.8.2 | | | | | |
| i> | 10.10.10.3/32 | | 12 | 25 | 0 | No | LSP_ | | |
| DEFAULT | 3 | xe2 | No | 10.1.5.2 | | | | | |
| i(b)> | 10.10.10.3/32 | | 8 | 21 | 2203 | Yes | LSP_ | | |
| DEFAULT | 17400 | xe2 | No | 10.1.5.2 | | | | | |
| i(b)> | 10.10.10.3/32 | | 9 | 22 | 2204 | Yes | LSP_ | | |
| DEFAULT | 17400 | xe0 | No | 10.1.8.2 | | | | | |
| i> | 10.10.10.5/32 | | 4 | 9 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe2 | No | 10.1.5.2 | | | | | |
| i> | 10.10.10.5/32 | | 13 | 27 | 0 | No | LSP_ | | |
| DEFAULT | 17600 | xe0 | No | 10.1.8.2 | | | | | |
| i> | 10.10.10.6/32 | | 5 | 15 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe0 | No | 10.1.8.2 | | | | | |
| i> | 10.10.10.6/32 | | 14 | 29 | 0 | No | LSP_ | | |
| DEFAULT | 17500 | xe2 | No | 10.1.5.2 | | | | | |

The following show output displays the FEC-To-NHLF (FTN) table information on PE1, PE2, PE3, and PE4 devices in the network [Figure 55](#) using the **show mpls ftn-table** command.

PE1#show mpls ftn-table

Primary FTN entry with FEC: 10.10.10.2/32, id: 1, row status: Active, Tunnel-Policy: N/A, State: Installed
 Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
 Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 12

```
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 12, owner: ISIS-SR, Stale: NO, out intf: xe2, out label: 17700
Nexthop addr: 10.1.2.2          cross connect ix: 7, op code: Push

Non-primary FTN entry with FEC: 10.10.10.2/32, id: 10, row status: Active, Tunnel-Policy: N/A,
State: Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xel, out label: 3
Nexthop addr: 10.1.1.2          cross connect ix: 2, op code: Push

bypass_ftn_ix 6

Primary FTN entry with FEC: 10.10.10.2/32, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 2201, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 16
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 16, owner: ISIS-SR, Stale: NO, out intf: xel, out label: 17700
Nexthop addr: 10.1.1.2          cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 10.10.10.3/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 13
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 13, owner: ISIS-SR, Stale: NO, out intf: xe2, out label: 17400
Nexthop addr: 10.1.2.2          cross connect ix: 3, op code: Push

Non-primary FTN entry with FEC: 10.10.10.3/32, id: 11, row status: Active, Tunnel-Policy: N/A,
State: Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xel, out label: 3
Nexthop addr: 10.1.1.2          cross connect ix: 2, op code: Push

bypass_ftn_ix 7

Primary FTN entry with FEC: 10.10.10.3/32, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 2202, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 18
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 18, owner: ISIS-SR, Stale: NO, out intf: xel, out label: 17400
Nexthop addr: 10.1.1.2          cross connect ix: 8, op code: Push

Primary FTN entry with FEC: 10.10.10.3/32, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 2204, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 13
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 13, owner: ISIS-SR, Stale: NO, out intf: xe2, out label: 17400
Nexthop addr: 10.1.2.2          cross connect ix: 3, op code: Push
```

Primary FTN entry with FEC: 10.10.10.4/32, id: 3, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 14

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 14, owner: ISIS-SR, Stale: NO, out intf: xe2, out label: 17300

Nexthop addr: 10.1.2.2 cross connect ix: 4, op code: Push

Non-primary FTN entry with FEC: 10.10.10.4/32, id: 12, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe1, out label: 3

Nexthop addr: 10.1.1.2 cross connect ix: 2, op code: Push

bypass_ftn_ix 8

Primary FTN entry with FEC: 10.10.10.4/32, id: 8, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 2203, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 20

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 20, owner: ISIS-SR, Stale: NO, out intf: xe1, out label: 17300

Nexthop addr: 10.1.1.2 cross connect ix: 9, op code: Push

Primary FTN entry with FEC: 10.10.10.5/32, id: 4, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe1, out label: 3

Nexthop addr: 10.1.1.2 cross connect ix: 2, op code: Push

Non-primary FTN entry with FEC: 10.10.10.5/32, id: 13, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 26

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 26, owner: ISIS-SR, Stale: NO, ISIS-SR out intf: xe2, transport out intf: N/A, out label: 17600

Nexthop addr: 10.1.2.2 cross connect ix: 10, op code: Push and Lookup

bypass_ftn_ix 9

Primary FTN entry with FEC: 10.10.10.6/32, id: 5, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 11

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 11, owner: N/A, Stale: NO, out intf: xe2, out label: 3

Nexthop addr: 10.1.2.2 cross connect ix: 6, op code: Push

Non-primary FTN entry with FEC: 10.10.10.6/32, id: 14, row status: Active, Tunnel-Policy: N/A, State: Installed

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
    Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 28
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 28, owner: ISIS-SR, Stale: NO, ISIS-SR out intf: xe1, transport out intf:
N/A, out label: 17500
    Nexthop addr: 10.1.1.2          cross connect ix: 11, op code: Push and Lookup

    bypass_ftn_ix 7

PE2#show mpls ftn-table
Primary FTN entry with FEC: 10.10.10.1/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 16
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 16, owner: ISIS-SR, Stale: NO, out intf: xe4, out label: 17800
    Nexthop addr: 10.1.3.2          cross connect ix: 7, op code: Push

Non-primary FTN entry with FEC: 10.10.10.1/32, id: 10, row status: Active, Tunnel-Policy: N/A,
State: Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
    Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe5, out label: 3
    Nexthop addr: 10.1.4.2          cross connect ix: 2, op code: Push

    bypass_ftn_ix 6

Primary FTN entry with FEC: 10.10.10.1/32, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 2201, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
    Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 20
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 20, owner: ISIS-SR, Stale: NO, out intf: xe5, out label: 17800
    Nexthop addr: 10.1.4.2          cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 10.10.10.3/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 17
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 17, owner: ISIS-SR, Stale: NO, out intf: xe4, out label: 17400
    Nexthop addr: 10.1.3.2          cross connect ix: 3, op code: Push

Non-primary FTN entry with FEC: 10.10.10.3/32, id: 11, row status: Active, Tunnel-Policy: N/A,
State: Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
    Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe5, out label: 3
    Nexthop addr: 10.1.4.2          cross connect ix: 2, op code: Push

    bypass_ftn_ix 7

Primary FTN entry with FEC: 10.10.10.3/32, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
```

```
Tunnel id: 2202, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
```

```
    Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 22
```

```
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
    Out-segment with ix: 22, owner: ISIS-SR, Stale: NO, out intf: xe5, out label: 17400
```

```
    Nexthop addr: 10.1.4.2          cross connect ix: 8, op code: Push
```

```
Primary FTN entry with FEC: 10.10.10.3/32, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
```

```
Tunnel id: 2204, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
```

```
    Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 17
```

```
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
    Out-segment with ix: 17, owner: ISIS-SR, Stale: NO, out intf: xe4, out label: 17400
```

```
    Nexthop addr: 10.1.3.2          cross connect ix: 3, op code: Push
```

```
Primary FTN entry with FEC: 10.10.10.4/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
    Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 18
```

```
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
    Out-segment with ix: 18, owner: ISIS-SR, Stale: NO, out intf: xe4, out label: 17300
```

```
    Nexthop addr: 10.1.3.2          cross connect ix: 4, op code: Push
```

```
Non-primary FTN entry with FEC: 10.10.10.4/32, id: 12, row status: Active, Tunnel-Policy: N/A,
State: Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
```

```
    Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
```

```
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
    Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe5, out label: 3
```

```
    Nexthop addr: 10.1.4.2          cross connect ix: 2, op code: Push
```

```
bypass_ftn_ix 8
```

```
Primary FTN entry with FEC: 10.10.10.4/32, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
```

```
Tunnel id: 2203, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
```

```
    Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 24
```

```
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
    Out-segment with ix: 24, owner: ISIS-SR, Stale: NO, out intf: xe5, out label: 17300
```

```
    Nexthop addr: 10.1.4.2          cross connect ix: 9, op code: Push
```

```
Primary FTN entry with FEC: 10.10.10.5/32, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
    Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
```

```
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
    Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe5, out label: 3
```

```
    Nexthop addr: 10.1.4.2          cross connect ix: 2, op code: Push
```

```
Non-primary FTN entry with FEC: 10.10.10.5/32, id: 13, row status: Active, Tunnel-Policy: N/A,
State: Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
```

```
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
```

```
    Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 30
```

```
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
```

```

    Out-segment with ix: 30, owner: ISIS-SR, Stale: NO, ISIS-SR out intf: xe4, transport out intf:
N/A, out label: 17600
    Nexthop addr: 10.1.3.2          cross connect ix: 10, op code: Push and Lookup

    bypass_ftn_ix 9

    Primary FTN entry with FEC: 10.10.10.6/32, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
    Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
    Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 15
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 15, owner: N/A, Stale: NO, out intf: xe4, out label: 3
    Nexthop addr: 10.1.3.2          cross connect ix: 6, op code: Push

    Non-primary FTN entry with FEC: 10.10.10.6/32, id: 14, row status: Active, Tunnel-Policy: N/A,
State: Installed
    Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
    Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
    Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 32
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 32, owner: ISIS-SR, Stale: NO, ISIS-SR out intf: xe5, transport out intf:
N/A, out label: 17500
    Nexthop addr: 10.1.4.2          cross connect ix: 11, op code: Push and Lookup

    bypass_ftn_ix 7

PE3#show mpls ftn-table
    Primary FTN entry with FEC: 10.10.10.1/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
    Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
    Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 12
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 12, owner: ISIS-SR, Stale: NO, out intf: xe5, out label: 17800
    Nexthop addr: 10.1.7.2          cross connect ix: 7, op code: Push

    Non-primary FTN entry with FEC: 10.10.10.1/32, id: 10, row status: Active, Tunnel-Policy: N/A,
State: Installed
    Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
    Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
    Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
    Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe1, out label: 3
    Nexthop addr: 10.1.6.2          cross connect ix: 2, op code: Push

    bypass_ftn_ix 6

    Primary FTN entry with FEC: 10.10.10.1/32, id: 6, row status: Active, Tunnel-Policy: N/A, State:
Installed
    Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
    Tunnel id: 2201, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
    Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 16
    Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
    Out-segment with ix: 16, owner: ISIS-SR, Stale: NO, out intf: xe1, out label: 17800
    Nexthop addr: 10.1.6.2          cross connect ix: 5, op code: Push

    Primary FTN entry with FEC: 10.10.10.1/32, id: 9, row status: Active, Tunnel-Policy: N/A, State:
Installed
    Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
    Tunnel id: 2204, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
    Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 12

```

```
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 12, owner: ISIS-SR, Stale: NO, out intf: xe5, out label: 17800
Nexthop addr: 10.1.7.2          cross connect ix: 7, op code: Push
```

```
Primary FTN entry with FEC: 10.10.10.2/32, id: 2, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 13
```

```
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 13, owner: ISIS-SR, Stale: NO, out intf: xe5, out label: 17700
```

```
Nexthop addr: 10.1.7.2          cross connect ix: 3, op code: Push
```

```
Non-primary FTN entry with FEC: 10.10.10.2/32, id: 11, row status: Active, Tunnel-Policy: N/A,
State: Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
```

```
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
```

```
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe1, out label: 3
```

```
Nexthop addr: 10.1.6.2          cross connect ix: 2, op code: Push
```

```
bypass_ftn_ix 7
```

```
Primary FTN entry with FEC: 10.10.10.2/32, id: 7, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 2202, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
```

```
Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 18
```

```
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 18, owner: ISIS-SR, Stale: NO, out intf: xe1, out label: 17700
```

```
Nexthop addr: 10.1.6.2          cross connect ix: 8, op code: Push
```

```
Primary FTN entry with FEC: 10.10.10.4/32, id: 3, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
```

```
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 14
```

```
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 14, owner: ISIS-SR, Stale: NO, out intf: xe5, out label: 17300
```

```
Nexthop addr: 10.1.7.2          cross connect ix: 4, op code: Push
```

```
Non-primary FTN entry with FEC: 10.10.10.4/32, id: 12, row status: Active, Tunnel-Policy: N/A,
State: Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
```

```
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
```

```
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe1, out label: 3
```

```
Nexthop addr: 10.1.6.2          cross connect ix: 2, op code: Push
```

```
bypass_ftn_ix 8
```

```
Primary FTN entry with FEC: 10.10.10.4/32, id: 8, row status: Active, Tunnel-Policy: N/A, State:
Installed
```

```
Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 2203, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, ,
Color: 0
```

```
Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 20
```

```
Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
```

```
Out-segment with ix: 20, owner: ISIS-SR, Stale: NO, out intf: xe1, out label: 17300
```

```
Nexthop addr: 10.1.6.2          cross connect ix: 9, op code: Push
```



```

Primary FTN entry with FEC: 10.10.10.5/32, id: 4, row status: Active, Tunnel-Policy: N/A, State:
Installed
  Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe1, out label: 3
  Nexthop addr: 10.1.6.2      cross connect ix: 2, op code: Push

Non-primary FTN entry with FEC: 10.10.10.5/32, id: 13, row status: Active, Tunnel-Policy: N/A,
State: Installed
  Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
  Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 26
  Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 26, owner: ISIS-SR, Stale: NO, ISIS-SR out intf: xe5, transport out intf:
N/A, out label: 17600
  Nexthop addr: 10.1.7.2      cross connect ix: 10, op code: Push and Lookup

bypass_ftn_ix 9

Primary FTN entry with FEC: 10.10.10.6/32, id: 5, row status: Active, Tunnel-Policy: N/A, State:
Installed
  Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 11
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 11, owner: N/A, Stale: NO, out intf: xe5, out label: 3
  Nexthop addr: 10.1.7.2      cross connect ix: 6, op code: Push

Non-primary FTN entry with FEC: 10.10.10.6/32, id: 14, row status: Active, Tunnel-Policy: N/A,
State: Installed
  Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
  Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 28
  Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 28, owner: ISIS-SR, Stale: NO, ISIS-SR out intf: xe1, transport out intf:
N/A, out label: 17500
  Nexthop addr: 10.1.6.2      cross connect ix: 11, op code: Push and Lookup

bypass_ftn_ix 7

PE4#show mpls ftn-table
Primary FTN entry with FEC: 10.10.10.1/32, id: 1, row status: Active, Tunnel-Policy: N/A, State:
Installed
  Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0
  Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 12
  Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 12, owner: ISIS-SR, Stale: NO, out intf: xe0, out label: 17800
  Nexthop addr: 10.1.8.2      cross connect ix: 7, op code: Push

Non-primary FTN entry with FEC: 10.10.10.1/32, id: 10, row status: Active, Tunnel-Policy: N/A,
State: Installed
  Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
  Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color:
0
  Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
  Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
  Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe2, out label: 3
  Nexthop addr: 10.1.5.2      cross connect ix: 2, op code: Push

bypass_ftn_ix 6

```

Primary FTN entry with FEC: 10.10.10.1/32, id: 6, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 2201, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 16

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 16, owner: ISIS-SR, Stale: NO, out intf: xe2, out label: 17800

Nexthop addr: 10.1.5.2 cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 10.10.10.2/32, id: 2, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 13

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 13, owner: ISIS-SR, Stale: NO, out intf: xe0, out label: 17700

Nexthop addr: 10.1.8.2 cross connect ix: 3, op code: Push

Non-primary FTN entry with FEC: 10.10.10.2/32, id: 11, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe2, out label: 3

Nexthop addr: 10.1.5.2 cross connect ix: 2, op code: Push

bypass_ftn_ix 7

Primary FTN entry with FEC: 10.10.10.2/32, id: 7, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 2202, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 18

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 18, owner: ISIS-SR, Stale: NO, out intf: xe2, out label: 17700

Nexthop addr: 10.1.5.2 cross connect ix: 8, op code: Push

Primary FTN entry with FEC: 10.10.10.3/32, id: 3, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 14

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 14, owner: ISIS-SR, Stale: NO, out intf: xe0, out label: 17400

Nexthop addr: 10.1.8.2 cross connect ix: 4, op code: Push

Non-primary FTN entry with FEC: 10.10.10.3/32, id: 12, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe2, out label: 3

Nexthop addr: 10.1.5.2 cross connect ix: 2, op code: Push

bypass_ftn_ix 8

Primary FTN entry with FEC: 10.10.10.3/32, id: 8, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none

Tunnel id: 2203, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 9, in intf: - in label: 0 out-segment ix: 20

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 20, owner: ISIS-SR, Stale: NO, out intf: xe2, out label: 17400

Nexthop addr: 10.1.5.2 cross connect ix: 9, op code: Push

Primary FTN entry with FEC: 10.10.10.3/32, id: 9, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none

Tunnel id: 2204, Protected LSP id: 0, LSP-type: Bypass, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 14

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 14, owner: ISIS-SR, Stale: NO, out intf: xe0, out label: 17400

Nexthop addr: 10.1.8.2 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 10.10.10.5/32, id: 4, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 2, owner: N/A, Stale: NO, out intf: xe2, out label: 3

Nexthop addr: 10.1.5.2 cross connect ix: 2, op code: Push

Non-primary FTN entry with FEC: 10.10.10.5/32, id: 13, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none

Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 26

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 26, owner: ISIS-SR, Stale: NO, ISIS-SR out intf: xe0, transport out intf: N/A, out label: 17600

Nexthop addr: 10.1.8.2 cross connect ix: 10, op code: Push and Lookup

bypass_ftn_ix 9

Primary FTN entry with FEC: 10.10.10.6/32, id: 5, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none

Tunnel id: 0, Protected LSP id: 0, LSP-type: Primary, Description: N/A, , Color: 0

Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 11

Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 11, owner: N/A, Stale: NO, out intf: xe0, out label: 3

Nexthop addr: 10.1.8.2 cross connect ix: 6, op code: Push

Non-primary FTN entry with FEC: 10.10.10.6/32, id: 14, row status: Active, Tunnel-Policy: N/A, State: Installed

Owner: ISIS-SR, distance: 115, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none

Tunnel id: 0, Protected LSP id: 0, LSP-type: Backup, QoS Resource id: 0, Description: N/A, , Color: 0

Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 28

Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up

Out-segment with ix: 28, owner: ISIS-SR, Stale: NO, ISIS-SR out intf: xe2, transport out intf: N/A, out label: 17500

Nexthop addr: 10.1.5.2 cross connect ix: 11, op code: Push and Lookup

bypass_ftn_ix 8

Port-Active

The following show output displays the Ethernet Segment (ES) and Intermediate System (IS) neighbor adjacencies for PE1, PE2, PE3, PE4, P1, and P2 devices in the network [Figure 55](#) using the **show clns neighbors** command.

```
PE1#show clns neighbors
```

```
Total number of L1 adjacencies: 2
```

```
Total number of L2 adjacencies: 2
```

```
Total number of adjacencies: 4
```

```
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P2 | xe2 | e8c5.7a55.3c7e | Up | 22 | L1 | IS-IS |
| | | | Up | 22 | L2 | IS-IS |
| P1 | xe1 | e49d.73b3.c107 | Up | 23 | L1 | IS-IS |

```
PE2#show clns neighbors
```

```
Total number of L1 adjacencies: 2
```

```
Total number of L2 adjacencies: 2
```

```
Total number of adjacencies: 4
```

```
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| P2 | xe4 | e8c5.7a55.3c7f | Up | 8 | L1 | IS-IS |
| | | | Up | 8 | L2 | IS-IS |
| P1 | xe5 | e49d.73b3.c14c | Up | 29 | L1 | IS-IS |
| | | | Up | 29 | L2 | IS-IS |
| | | | Up | 23 | L2 | IS-IS |

```
P1#show clns neighbors
```

```
Total number of L1 adjacencies: 4
```

```
Total number of L2 adjacencies: 4
```

```
Total number of adjacencies: 8
```

```
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| PE3 | xe4 | b86a.97d9.2cdf | Up | 19 | L1 | IS-IS |
| | | | Up | 19 | L2 | IS-IS |
| PE2 | xe2 | e8c5.7a47.9dfc | Up | 7 | L1 | IS-IS |
| | | | Up | 7 | L2 | IS-IS |
| PE1 | xe1 | e8c5.7a78.c918 | Up | 7 | L1 | IS-IS |
| | | | Up | 7 | L2 | IS-IS |
| PE4 | xe3 | d077.ceda.7004 | Up | 19 | L1 | IS-IS |
| | | | Up | 19 | L2 | IS-IS |

```
P2#show clns neighbors
```

```
Total number of L1 adjacencies: 4
```

```
Total number of L2 adjacencies: 4
```

```
Total number of adjacencies: 8
```

```
Tag 1: VRF : default
```

| System Id | Interface | SNPA | State | Holdtime | Type | Protocol |
|-----------|-----------|----------------|-------|----------|------|----------|
| PE3 | xe11 | b86a.97d9.2ccb | Up | 19 | L1 | IS-IS |
| | | | Up | 19 | L2 | IS-IS |
| PE1 | xe12 | e8c5.7a78.c908 | Up | 7 | L1 | IS-IS |
| | | | Up | 7 | L2 | IS-IS |
| PE2 | xe13 | e8c5.7a47.9dfb | Up | 19 | L1 | IS-IS |
| | | | Up | 19 | L2 | IS-IS |
| PE4 | xe14 | d077.ceda.7002 | Up | 19 | L1 | IS-IS |
| | | | Up | 19 | L2 | IS-IS |

```
PE3#show clns neighbors
```

```
Total number of L1 adjacencies: 2
```

```
Total number of L2 adjacencies: 2
```

```
Total number of adjacencies: 4
Tag 1: VRF : default
System Id      Interface      SNPA              State Holdtime  Type Protocol
P2              xe5              e8c5.7a55.3c77    Up    7          L1    IS-IS
                xe5              e8c5.7a55.3c77    Up    7          L2    IS-IS
P1              xe1              e49d.73b3.c105    Up    5          L1    IS-IS
                xe1              e49d.73b3.c105    Up    5          L2    IS-IS
```

```
PE4#show clns neighbors
```

```
Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 2
Total number of adjacencies: 4
Tag 1: VRF : default
System Id      Interface      SNPA              State Holdtime  Type Protocol
P2              xe0              e8c5.7a55.3c80    Up    5          L1    IS-IS
                xe0              e8c5.7a55.3c80    Up    5          L2    IS-IS
P1              xe2              e49d.73b3.c14d    Up    6          L1    IS-IS
                xe2              e49d.73b3.c14d    Up    6          L2    IS-IS
```

Port-Active ELAN

The following show outputs provide validation for ELAN configurations.

The following show output displays the EVPN active multi-homed and load-balanced details on PE1, PE2, PE3, and PE4 devices in the network [Figure 55](#) using the show evpn load-balance port-active and show evpn multi-homing all commands.

```
PE1#show evpn load-balance port-active
ESI              AC-IF/PE      PE-IP-ADDRESS    Redundancy      Service-
carving  weight  Revertive  AC-DF  Status
=====
00:00:00:11:11:77:77:00:00:00  LOCAL      10.10.10.1      port-
active      auto      0      NO      NA      STANDBY
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.2      port-
active      auto      0      NO      NA      ACTIVE
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.3      port-active      ----      --
--      ----      ----      STANDBY
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.4      port-active      ----      --
--      ----      ----      ACTIVE

PE2#show evpn load-balance port-active
ESI              AC-IF/PE      PE-IP-ADDRESS    Redundancy      Service-
carving  weight  Revertive  AC-DF  Status
=====
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.1      port-
active      auto      0      NO      NA      STANDBY
00:00:00:11:11:77:77:00:00:00  LOCAL      10.10.10.2      port-
active      auto      0      NO      NA      ACTIVE
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.3      port-active      ----      --
--      ----      ----      STANDBY
00:00:00:22:22:77:77:00:00:00  REMOTE      10.10.10.4      port-active      ----      --
--      ----      ----      ACTIVE

PE3#show evpn load-balance port-active
ESI              AC-IF/PE      PE-IP-ADDRESS    Redundancy      Service-
carving  weight  Revertive  AC-DF  Status
=====
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.1      port-active      ----      --
--      ----      ----      STANDBY
00:00:00:11:11:77:77:00:00:00  REMOTE      10.10.10.2      port-active      ----      --
--      ----      ----      ACTIVE
00:00:00:22:22:77:77:00:00:00  LOCAL      10.10.10.3      port-
active      auto      0      NO      NA      STANDBY
```

```

00:00:00:22:22:77:77:00:00:00 REMOTE 10.10.10.4 port-
active auto 0 NO NA ACTIVE

PE4#show evpn load-balance port-active
ESI AC-IF/PE PE-IP-ADDRESS Redundancy Service-
carving weight Revertive AC-DF Status
=====
00:00:00:11:11:77:77:00:00:00 REMOTE 10.10.10.1 port-active ---- --
-- ---- STANDBY
00:00:00:11:11:77:77:00:00:00 REMOTE 10.10.10.2 port-active ---- --
-- ---- ACTIVE
00:00:00:22:22:77:77:00:00:00 REMOTE 10.10.10.3 port-
active auto 0 NO NA STANDBY
00:00:00:22:22:77:77:00:00:00 LOCAL 10.10.10.4 port-
active auto 0 NO NA ACTIVE

PE1#show evpn multi-homing all
ESI Access-IF PE-IP-ADDRESS
=====
00:00:00:11:11:77:77:00:00:00 po1 10.10.10.1
00:00:00:11:11:77:77:00:00:00 ---- 10.10.10.2
00:00:00:22:22:77:77:00:00:00 ---- 10.10.10.3
00:00:00:22:22:77:77:00:00:00 ---- 10.10.10.4
Total number of entries are 4

PE2#show evpn multi-homing all
ESI Access-IF PE-IP-ADDRESS
=====
00:00:00:11:11:77:77:00:00:00 ---- 10.10.10.1
00:00:00:11:11:77:77:00:00:00 po1 10.10.10.2
00:00:00:22:22:77:77:00:00:00 ---- 10.10.10.3
00:00:00:22:22:77:77:00:00:00 ---- 10.10.10.4
Total number of entries are 4

PE3#show evpn multi-homing all
ESI Access-IF PE-IP-ADDRESS
=====
00:00:00:11:11:77:77:00:00:00 ---- 10.10.10.1
00:00:00:11:11:77:77:00:00:00 ---- 10.10.10.2
00:00:00:22:22:77:77:00:00:00 po1 10.10.10.3
00:00:00:22:22:77:77:00:00:00 ---- 10.10.10.4
Total number of entries are 4

PE4#show evpn multi-homing all
ESI Access-IF PE-IP-ADDRESS
=====
00:00:00:11:11:77:77:00:00:00 ---- 10.10.10.1
00:00:00:11:11:77:77:00:00:00 ---- 10.10.10.2
00:00:00:22:22:77:77:00:00:00 ---- 10.10.10.3
00:00:00:22:22:77:77:00:00:00 po1 10.10.10.4
Total number of entries are 4

```

The following show output displays the active EVPN MPLS Tunnels for ELAN on PE1, PE2, PE3, and PE4 devices in the network [Figure 55](#) using the show evpn mpls tunnel command.

```

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source Destination Status Up/Down Update evpn-id
=====
10.10.10.1 10.10.10.2 Installed 00:02:19 00:02:19 3000
10.10.10.1 10.10.10.4 Installed 00:10:21 00:10:21 3000
10.10.10.1 10.10.10.3 Installed 00:10:47 00:10:47 3000

Total number of entries are 3

PE2#show evpn mpls tunnel

```

```

EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      evpn-id
=====
10.10.10.2      10.10.10.1      Installed   00:02:01     00:02:01     3000
10.10.10.2      10.10.10.4      Installed   00:02:01     00:02:01     3000
10.10.10.2      10.10.10.3      Installed   00:02:01     00:02:01     3000

Total number of entries are 3

PE3#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      evpn-id
=====
10.10.10.3      10.10.10.2      Installed   00:02:27     00:02:27     3000
10.10.10.3      10.10.10.4      Installed   00:10:29     00:10:29     3000
10.10.10.3      10.10.10.1      Installed   00:10:54     00:10:54     3000

Total number of entries are 3

PE4#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      evpn
-id
=====
===
10.10.10.4      10.10.10.2      Installed   00:02:30     00:02:30     3000
10.10.10.4      10.10.10.3      Installed   00:10:32     00:10:32     3000
10.10.10.4      10.10.10.1      Installed   00:10:32     00:10:32     3000

Total number of entries are 3

```

Port-Active ELINE

The following show output displays the active EVPN SR Tunnels for ELINE on PE1, PE2, PE3, and PE4 devices in the network [Figure 55](#) using the `show evpn mpls xconnect tunnel` command.

```

PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====
===
10.10.10.1      10.10.10.4      AC-Down     01:07:01     01:07:01     1800          1700
10.10.10.1      10.10.10.3      AC-Down     01:07:01     01:07:01     1800          1700

Total number of entries are 2

PE2#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====
===
10.10.10.2      10.10.10.3      Installed   00:08:20     00:07:31     1800          1700
10.10.10.2      10.10.10.4      Installed   00:08:20     00:07:31     1800          1700

Total number of entries are 2

PE3#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id
=====

```

```

===
10.10.10.3      10.10.10.1      AC-Down      01:04:48      01:04:48      1700      1800
10.10.10.3      10.10.10.2      AC-Down      00:08:48      00:08:48      1700      1800

Total number of entries are 2

PE4#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status          Up/Down          Update          local-evpn-id remote-
evpn-id
=====
===
10.10.10.4      10.10.10.1      Installed      00:09:00      00:08:28      1700      1800
10.10.10.4      10.10.10.2      Installed      00:09:01      00:08:28      1700      1800

Total number of entries are 2

```

Implementation Examples

Scenario: Customer wants to achieve redundancy for its hosts in a network using Single-Active or Port-Active redundancy.

- Customer configures the `evpn multi-homed` command with the `load-balancing single-active` or `load-balancing port-active` option on the relevant PE interfaces.
- Single-Active or Port-Active redundancy is now in effect, ensuring redundancy for hosts.
- The feature works in conjunction with other EVPN-related configurations, such as VRF, VLAN mapping, and other EVPN settings.

Revised CLI Commands

Below is the revised command for configuring EVPN Active-Standby.

evpn multi-homed

- The command `evpn multi-homed` allows users to configure single-active and port-active load-balancing Ethernet Segment Identifier (ESI) on a link with a multihomed Customer Edge (CE) in the context of EVPN multi-homed configurations. For more details, refer to the [evpn multi-homed \(page 1563\)](#) command in the [EVPN MPLS Commands \(page 1538\)](#) chapter in the Multi-Protocol Label Switching Guide.
- The existing syntax now includes the newly added parameter for load-balancing, namely `single-active` and `port-active`.

CLI Commands

The EVPN Active-Standby feature introduces the following configuration commands. For more information of the EVPN MPLS commands, see the [EVPN MPLS Commands \(page 1538\)](#) chapter in the Multi-Protocol Label Switching Guide.

service-carving ac-driven

Use this command to enable the AC-influenced method for any selected Designated Forwarder (DF) algorithm.

Enabling the `ac-driven` method allows the Designated Forwarder (DF) algorithm to be influenced by the Attachment Circuits (ACs) associated with a specific Ethernet Segment (ES). This means that the DF selection is based on the ACs characteristics and conditions, such as whether an AC is operational UP, mapped, or unmapped on the ESI.

Use `no` form of this command to disable the AC-influenced method for any selected Designated Forwarder (DF) algorithm.

Command Syntax

```
service-carving ac-driven
no service-carving ac-driven
```

Parameters

None

Default

`ac-driven` is disabled.

Command Mode

EVPN Ethernet Segment (ES) Mode

Applicability

This command was introduced in the OcNOS version 6.4.2.

Example

The provided examples showcase the configuration of the `service-carving ac-driven` command in EVPN Ethernet Segment (ES) mode. The first two examples demonstrate enabling this feature with different DF election methods, and the final example illustrates the command to disable `service-carving ac-driven`.

```
OcNOS#configure terminal
OcNOS(config)#interface sal
OcNOS(config-if)#evpn multi-homed esi 11:22:33:44:55:66:77:88:99 load-balancing single-active
OcNOS(config-if-es)#service-carving preference-based
OcNOS(config-if-es)#service-carving ac-driven
OcNOS(config-if-es)#end

OcNOS#configure terminal
OcNOS(config)#interface sal
OcNOS(config-if)#evpn multi-homed esi 11:22:33:44:55:66:77:88:99 load-balancing single-active
OcNOS(config-if-es)#service-carving auto
OcNOS(config-if-es)#service-carving ac-driven

OcNOS(config-if-es)#no service-carving ac-driven
OcNOS(config-if-es)#end
```

service-carving

Use this command to provide the flexibility to select the Designated Forwarder (DF) election algorithm based on preference based or modulo-based DF election.

Use no form of this command to disable service-carving.

Command Syntax

```
service-carving (preference-based|auto)
no service-carving
```

Parameters

preference-based

Select the DF election algorithm based on preference based.

auto

Select the DF election algorithm based on modulo based.

Default

None

Command Mode

EVPN ES Mode

Applicability

This command was introduced in the OcNOS version 6.4.1.

Example

The following examples demonstrate the configuration of the `service-carving` command in both `single-active` or `port-active` mode for the EVPN multi-homed system, with one utilizing `auto` service carving and the other using `preference-based` service carving.

```
OcNOS#configure terminal
OcNOS(config)#interface sal
OcNOS(config-if)#evpn multi-homed esi 11:22:33:44:55:66:77:88:99 load-balancing single-active
OcNOS(config-if-es)#service-carving auto
OcNOS(config-if-es)#end

OcNOS#configure terminal
OcNOS(config)#interface pol
OcNOS(config-if)#evpn multi-homed system-mac 0000.0000.0011 load-balancing port-active
OcNOS(config-if-es)#service-carving auto
OcNOS(config-if-es)#end

OcNOS#configure terminal
OcNOS(config)#interface pol
OcNOS(config-if)#evpn multi-homed system-mac 0000.0000.0011 load-balancing port-active
OcNOS(config-if-es)#service-carving preference-based
OcNOS(config-if-es)#end
```

The following example is used to disable the `service-carving` for the EVPN multi-homed system.

```
OcNOS(config-if-es)#no service-carving
OcNOS(config-if-es)#end
```

service-carving weight

Use this command to specify a preference value when the preference-based Designated Forwarder (DF) election algorithm is selected. This preference value determines the priority of the local PE device to become the DF for a particular Ethernet segment.

Use no form of this command to replace the preference weight value and choose the default preference value.

Command Syntax

```
service-carving weight <1-65535>  
no service-carving weight
```

Parameters

weight <1-65535>

Specifies the preference weight value. A lower weight value indicates a higher priority for becoming the DF.

Default

The service-carving weight command is set to 32767 by default.

Command Mode

EVPN Ethernet Segment (ES) Mode

Applicability

This command was introduced in the OcNOS version 6.4.1.

Example

The `service-carving weight` command is used to configure the preference weight value for service-carving in both port-active and single-active modes.

```
OcNOS#configure terminal  
OcNOS(config)#interface pol  
OcNOS(config-if)#evpn multi-homed system-mac 0000.0000.0011 load-balancing port-active  
OcNOS(config-if-es)#service-carving preference-based  
OcNOS(config-if-es)#service-carving weight 100  
OcNOS(config-if-es)#end  
  
OcNOS#configure terminal  
OcNOS(config)#interface sal  
OcNOS(config-if)#evpn multi-homed esi 11:22:33:44:55:66:77:88:99 load-balancing single-active  
OcNOS(config-if-es)#service-carving preference-based  
OcNOS(config-if-es)#service-carving weight 100
```

To disable the configured weight, use the `no service-carving weight` command.

```
OcNOS(config-if-es)#no service-carving weight  
OcNOS(config-if-es)#end
```

Troubleshooting

To ensure the reliable operation of the single-active or port-active setup and maintain data accuracy and consistency, follow these troubleshooting steps:

1. Verify the Configuration:

- Use the `show running-config` command to confirm that the ESI configuration includes load-balancing single-active or port-active, such as:

```
evpn multi-homed esi 11:22:33:44:55:66:77:88:99 load-balancing single-active
or
evpn multi-homed system-mac 0000.4321.1234 load-balancing port-active
```

- Ensure that the `service-carving` algorithm type is configured.

2. Verify the `show` command:

- Use the `show bgp l2vpn evpn multihoming es-route` command to confirm that it matches the `service-carving` algorithm type.
- Use the `show evpn load-balance single-active` or `port-active` command to verify the status of the Multihomed (MH) nodes as `ACTIVE` and `STANDBY`.

3. Ensure Proper Connectivity:

Validate the connectivity between the router and the EVPN tunnel to ensure it is up. This involves verifying network settings, ports, and firewalls.

4. For the server:

Enable debugging on OcNOS and enable debug mode. Verify the logs in `/var/log/messages` for further insights.

Glossary

The following provides definitions for key terms used throughout this document.

| Key Terms or Acronym | Description |
|---|--|
| EVPN | Ethernet Virtual Private Network |
| ELINE | Ethernet Line services |
| ELAN | Ethernet LAN services |
| LAN | Local Area Network |
| CE | Customer Edge |
| PE | Provider Edge |
| MH | Multihoming |
| AC | Attachment Circuit |
| LACP | Link Aggregation Control Protocol |
| BUM | Broadcast, Unknown Unicast, Multicast |
| MAC | Media Access Control |
| ARP/ND | Address Resolution Protocol/Neighbor Discovery |
| DF | Designated Forwarder |
| Ethernet Virtual Private Network (EVPN) | A network technology that extends Layer 2 Ethernet services over a Layer 3 IP/MPLS network. |
| Ethernet Line services (ELINE) | Two PEs are directly connected over an Ethernet link, enabling redundancy and efficient data exchange. |
| Ethernet LAN | A group of PEs are interconnected in a multipoint Ethernet network, providing redundancy |

| Key Terms or Acronym | Description |
|---|---|
| services (ELAN) | and optimized data transfer. |
| Port-Active | A redundancy mechanism in which multiple Provider Edge (PE) devices can be active simultaneously for the same host or MAC address, with specific active ports associated with each active PE. |
| Single-Active | A redundancy mechanism in which only one of the Provider Edge (PE) devices is active at a time for handling traffic for a specific host or MAC address. |
| Customer Edge (CE) | A device at the customer's network edge that connects to the service provider's network. |
| Provider Edge (PE) | A device at the service provider's network edge that connects to customer edge devices. |
| Multihoming (MH) | Connecting a host or CE device to multiple PE devices for redundancy and load balancing. |
| Attachment Circuit (AC) | The connection between a CE device and a PE device in an EVPN network. |
| Link Aggregation Control Protocol (LACP) | A protocol used to manage and bundle multiple physical links into a single logical link for higher bandwidth and redundancy. |
| Broadcast, Unknown Unicast, Multicast (BUM) | Categories of network traffic that includes broadcast, unknown unicast, and multicast packets. |
| Media Access Control (MAC) | A unique identifier assigned to network interfaces, typically associated with a hardware address. |
| Address Resolution Protocol/Neighbor Discovery (ARP/ND) | Protocols used to map IP addresses to MAC addresses in a local network. |
| Designated Forwarder (DF) | A PE device selected to forward broadcast, unknown unicast, and multicast traffic within an Ethernet segment. |
| Redundancy | The provision of duplicate equipment or links to ensure network availability in case of failures. |
| Failover | The process of switching to a backup device or link in case of a primary device or link failure. |
| Resiliency | The ability of a network to maintain its functionality even in the face of failures or disruptions. |
| Unicast | Communication between a single sender and a single receiver in a network. |
| Multicast | Communication from a single sender to multiple receivers in a network. |
| Egress | The process of traffic leaving a device or network segment. |
| Standby | In redundancy, a secondary device or link that is ready to take over in case the primary device or link fails. |
| Active | In redundancy, the primary device or link that is currently handling traffic. |
| Forwarding | The process of transmitting network packets from one device to another. |

| Key Terms or Acronym | Description |
|--------------------------------------|---|
| Link State | The operational status of a network link, indicating whether it is up or down. |
| Virtual Routing and Forwarding (VRF) | A technology that enables multiple instances of a routing table to coexist within a router. |
| Virtual Local Area Network (VLAN) | A logical network segment within a physical network. |
| Data Exchange | The process of sending and receiving data between network devices. |
| Downtime | The period during which a network or service is not available due to maintenance or failures. |

EVPN MPLS E-Tree Scenario 1

Overview

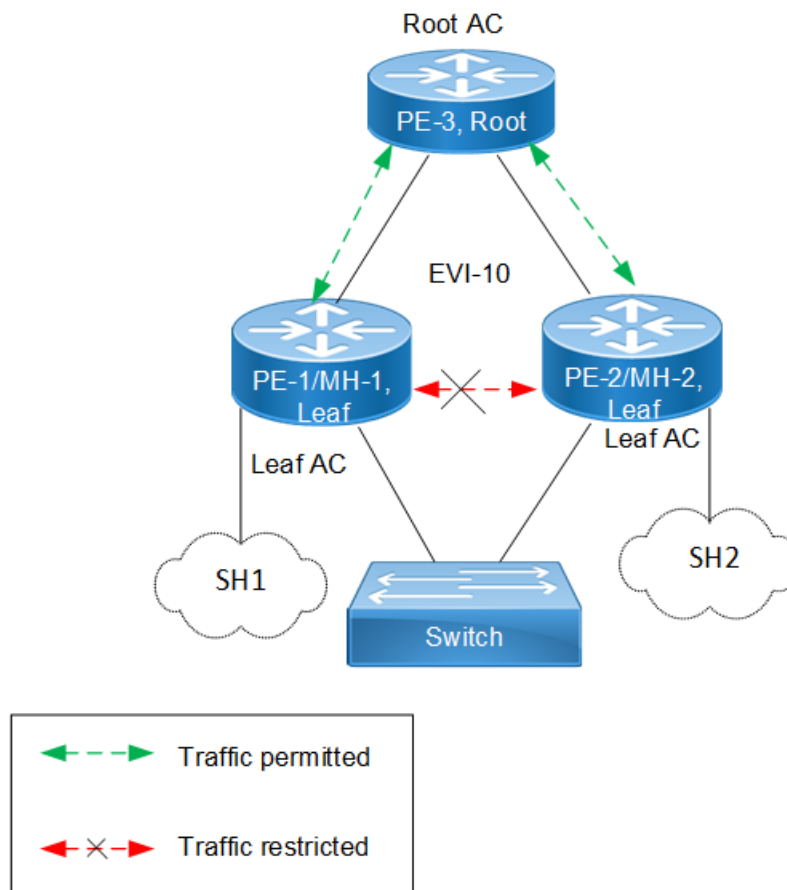
Ethernet VPN Ethernet-Tree (EVPN E-Tree), is a networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, especially in scenarios where specific services or devices need controlled communication. It categorizes network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them. OcNOS supports EVPN MPLS E-Tree based on Scenario 1 (SC-1) of RFC 8317, designating each Provider Edge (PE) node as either a Leaf or a Root site for Virtual Private Network (VPN) using VXLAN and MPLS EVPN.

Feature Characteristics

Scenario 1: Leaf or Root Site(s) per PE

Scenario 1 involves a topology with three PE nodes: PE-1, PE-2, and PE-3. PE-1 and PE-2 are Multi-Homed nodes (MH-1 and MH-2), with PE-3 acting as the Root node. PE-1 and PE-2 function as Leaf nodes and are part of a single home access interface (SH1 and SH2).

Figure 56. EVPN E-Tree



The classification ensures that communication follows specific rules:

- Communication between Leaf hosts is restricted, as indicated by red dotted lines with a cross mark (X) in the topology diagram. However, communication between Leaf and Root nodes, as well as between Root nodes, is permitted, marked by green dotted lines.
- Leaf nodes within PE-1 and PE-2 are isolated from each other, preventing intra-PE communication.

The scenario 1 is achieved through two main concepts:

1. Inter-PE Communication

- The inter-PE Route Target (RT) Constraint Method is applicable only to Single-Homing (SH) devices. Two RTs per broadcast domain are utilized, with Leaf PEs exporting Leaf RTs and Root nodes exporting Root RTs. Leaf nodes import only Root RTs, allowing communication with Root PEs while preventing communication with other Leaf nodes. RT constraints limit the import of specific EVPN routes (MAC-IP and IMET routes) to designated paths for inter-PE communication.
- IPI employs a proprietary method to support inter-PE connectivity for both SH and MH devices, using BGP extended community to advertise Leaf Indication in BGP routes and influence traffic flow for both Unicast and BUM traffic. This method enables implementation of ARP or ND cache suppression and MAC mobility sub-features specified in RFC-7432.

2. Intra-PE communication: Local Split Horizon controls intra-PE communication between Attachment Circuits (ACs) within Leaf PE nodes, ensuring that traffic between ACs does not egress to other Leaf ACs.



Note: This functionality depends on hardware capabilities.

Benefits

EVPN E-Tree Scenario 1 offers benefits in networking environments by providing efficient traffic control, enhanced security, scalability, and improved performance.

Efficient Traffic Control: Controls traffic between Leaf and Root sites across PEs, ensuring only authorized communication flows.

Enhanced Security: Leaf-to-Leaf traffic is blocked and split-horizon prevents local Leaf-to-Leaf communication, reducing risk of unauthorized access.

Scalability: Supports multiple Leaf and Root sites across the network, scalable for small to large enterprise deployments, but each PE hosts either Leaf or Root.

Improved Performance: Optimizes forwarding and filtering paths within SC-1 deployment, ensuring efficient delivery of critical data.

Prerequisites

In setting up a MPLS EVPN network, certain prerequisites are essential to ensure proper functionality and connectivity.

- **Ensure MPLS EVPN Configuration:** Confirm that MPLS EVPN and MPLS MH filtering are already enabled in all leaf and root nodes of the network as they are required for MPLS EVPN Multihoming.

```
!hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
!
evpn mpls multihoming enable
!
qos enable
!
```


- **Define Interfaces and Loopback Addresses:** Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific system MAC addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

```

!
interface po1
  switchport
  load-interval 30
  evpn multi-homed system-mac 0000.4321.1234
!
interface lo
  ip address 8.8.8.8/32 secondary
  ip router isis ISIS-IGP
!
interface xe8
  switchport
!
interface xe26
  channel-group 1 mode active
!

```

- **Configure ISIS and BGP for Dynamic Routing:** Enable ISIS to facilitate dynamic routing on all Leaf and Root nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Additionally, establish BGP sessions to advertise routes between different nodes. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

```

!
router isis ISIS-IGP
  is-type level-1
  ignore-lsp-errors
  lsp-gen-interval 5
  spf-interval-exp level-1 50 2000
  metric-style wide
  mpls traffic-eng router-id 8.8.8.8
  mpls traffic-eng level-1
  capability cspf
  dynamic-hostname
  fast-reroute terminate-hold-on interval 10000
  fast-reroute per-prefix level-1 proto ipv4 all
  fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
  bfd all-interfaces
  net 49.0001.0000.0000.0008.00
!
router bgp 65535
  neighbor 9.9.9.9 remote-as 65535
  neighbor 24.24.24.24 remote-as 65535
  neighbor 26.26.26.26 remote-as 65535
  neighbor 29.29.29.29 remote-as 65535
  neighbor 9.9.9.9 update-source lo
  neighbor 9.9.9.9 fall-over bfd
  neighbor 24.24.24.24 update-source lo
  neighbor 24.24.24.24 fall-over bfd
  neighbor 26.26.26.26 update-source lo
  neighbor 26.26.26.26 fall-over bfd
  neighbor 29.29.29.29 update-source lo
  neighbor 29.29.29.29 fall-over bfd
!
  address-family l2vpn evpn
  neighbor 9.9.9.9 activate
  neighbor 24.24.24.24 activate
  neighbor 26.26.26.26 activate
  neighbor 29.29.29.29 activate
  exit-address-family

```

```
!
exit
!
```

- **Configure LDP and RSVP for Efficient Network Operation:** Enable Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP) on all Leaf and Root nodes to optimize traffic routing and quality of service. LDP assigns labels for packet forwarding, while RSVP reserves network resources along specified paths to enhance network performance and reliability.

```
!
router ldp
  router-id 8.8.8.8
  fast-reroute
  graceful-restart full
  graceful-restart timers neighbor-liveness 120
  graceful-restart timers max-recovery 120
  session-protection duration 10
  targeted-peer ipv4 9.9.9.9
    exit-targeted-peer-mode
  targeted-peer ipv4 24.24.24.24
    exit-targeted-peer-mode
  transport-address ipv4 8.8.8.8
!
router rsvp
!
rsvp-path LEAF1-ROOT2 mpls
  24.1.4.24 strict
!
rsvp-path LEAF1-ROOT1 mpls
  26.1.2.26 strict
!
rsvp-trunk LEAF1-ROOT1 ipv4
  primary fast-reroute protection facility
  primary path LEAF1-ROOT1
  to 9.9.9.9
!
rsvp-trunk LEAF1-ROOT2 ipv4
  primary fast-reroute protection facility
  primary path LEAF1-ROOT2
  to 24.24.24.24
!
```

- **Create VRF for Isolated Routing Instances:** Configure VRF on all Leaf and Root nodes to create isolated routing instances within the network. This enables separate routing tables and forwarding behaviors for different groups of network resources.

```
!
mac vrf vrf103
  rd 8.8.8.8:103
  route-target both 65535:103
!
```

- **Connect Network Interfaces:** Configure network interfaces on all Leaf and Root nodes with connection details, IP addresses, and protocol settings. Enable label-switching and configure participation in the ISIS routing protocol, including support for protocols like LDP and RSVP for IPv4. These configurations optimize routing and resource management across the network.

```
!
interface xe11
  description connected to ROOT2 int xe9
  ip address 24.1.4.25/24
  label-switching
  ip router isis ISIS-IGP
  enable-ldp ipv4
  enable-rsvp
```

```

!
interface xe20
description connected to ROOT1 int xe20
ip address 26.1.2.27/24
label-switching
ip router isis ISIS-IGP
enable-ldp ipv4
enable-rsvp
!

```

- **Configure Switch:** Set up a VLAN bridge by enabling the VLAN and associating specific VLANs with the bridge. Configure network interfaces as trunk ports to allow traffic for all permitted VLANs across the network. Designate interfaces connected to Leaf and Root nodes as member ports of the VLAN bridge. This setup optimizes network segmentation and traffic management

```

!
bridge 1 protocol rstp vlan-bridge
!
vlan database
vlan-reservation 4030-4094
vlan 2-3010 bridge 1 state enable
!
interface po100
switchport
bridge-group 1
switchport mode trunk
switchport trunk allowed vlan all
!
interface lo
ip address 32.32.32.32/32 secondary
!
interface xe9
channel-group 100 mode active
!
interface xe17
channel-group 100 mode active
!
interface xe1
switchport
bridge-group 1
switchport mode trunk
switchport trunk allowed vlan all
!
exit
!

```

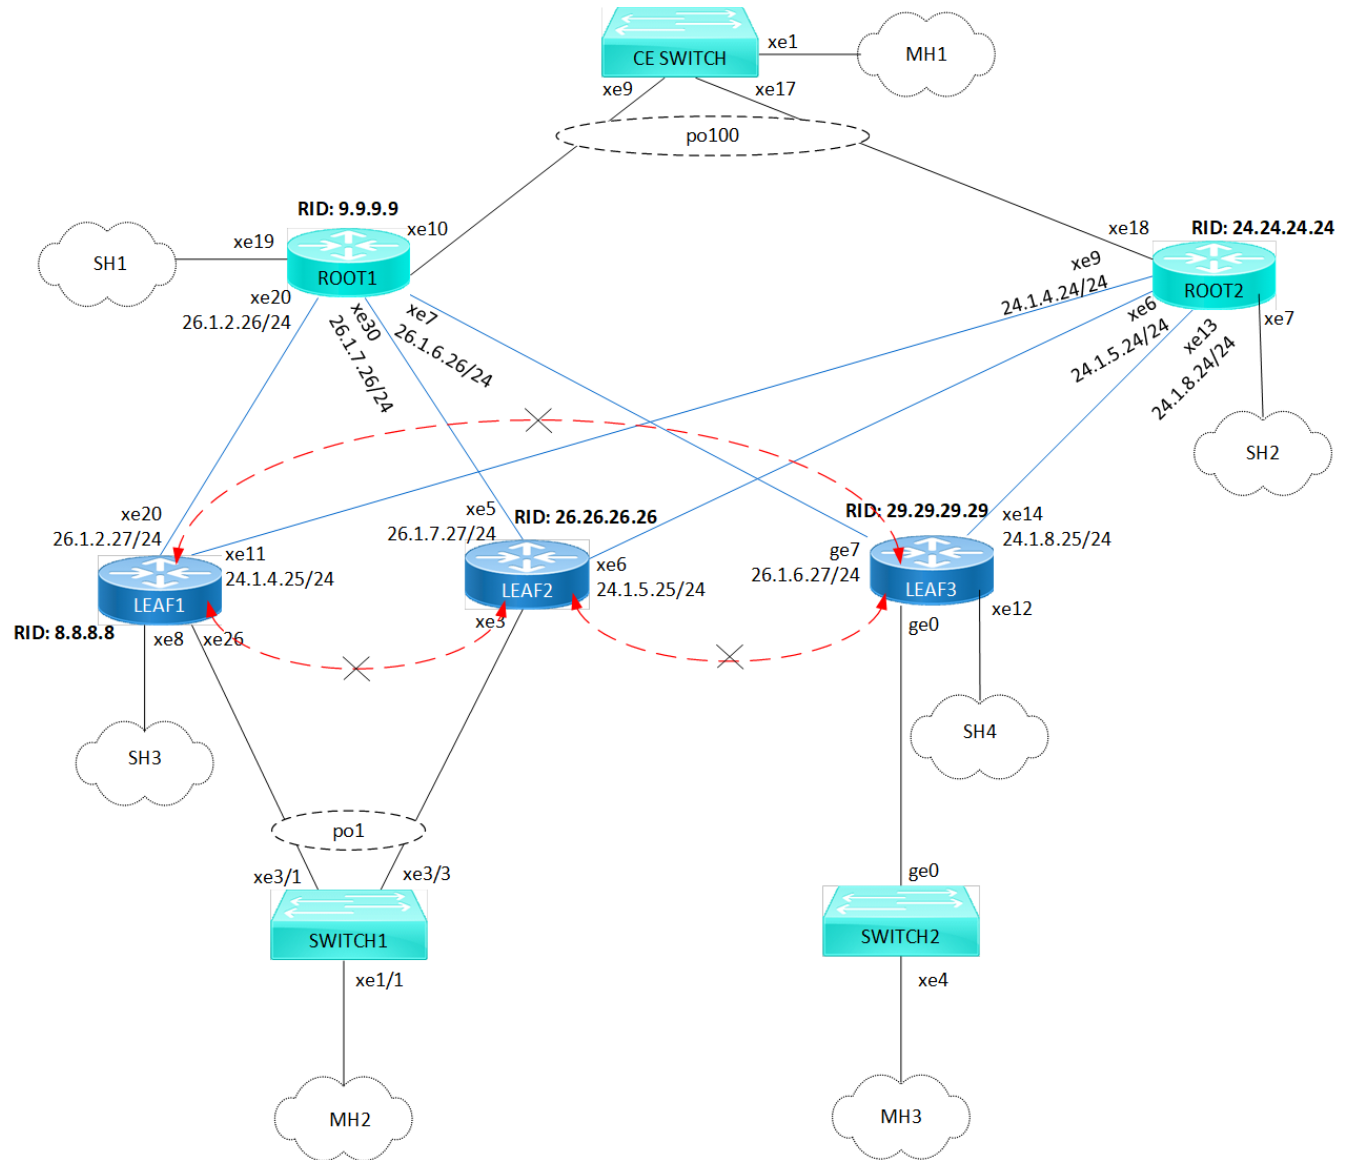
Configuration

Configure various nodes within the topology to set up an MPLS EVPN E-Tree Scenario 1 network, ensuring EVPN E-Tree for All-Active and Active-Standby redundancy and load balancing.

Topology

In the sample topology, Leaf nodes (LEAF1, LEAF2, LEAF3, and LEAF4), Root nodes (ROOT1 and ROOT2), and Switches (CE SWITCH, SWITCH1, and SWITCH2) form the network architecture. LEAF1 and LEAF2 are part of a Multi-homed group, with both connected to po1 (MH2). LEAF1 and LEAF3 have single home access-if ports (SH3 and SH4, respectively). Similarly, ROOT1 and ROOT2 are part of a Multi-homed group with po100 (MH1), and they each have a single home access-if port (SH1 and SH2, respectively). Leaf nodes are interconnected, and CE SWITCH, SWITCH1, and SWITCH2 are configured for Multi-homed connections to Leaf and Root nodes. SWITCH1 connects to LEAF1 and LEAF2, while CE SWITCH links to ROOT1 and ROOT2.

Figure 57. MPLS EVPN E-Tree Topology



BGP ID: 65535
 ISIS Instance: ISIS-IGP
 Leaf nodes VNID:203
 EVPN MH System MAC

- po1: 0000.4321.1234
- po100: 0000.1111.2222

← X → Traffic between leaf nodes is restricted.



Note: Before configuring E-Tree SC-1, meet all [Prerequisites \(page 1860\)](#) for the following nodes:

- Leaf nodes: LEAF1, LEAF2, and LEAF3
- Root nodes: ROOT1 and ROOT2
- Switches: CE SWITCH, SWITCH1 and SWITCH2

Enable EVPN E-Tree

The following E-Tree SC-1 configurations applies to Leaf and Root nodes within the MPLS network.

1. Enable EVPN E-Tree SC-1 which allows the nodes to participate in E-Tree functionality within the network, controlling traffic and establishing hierarchical connections between Leaf nodes in the network architecture.

```
(config)#evpn etree enable
```

2. Set the MAC ageing time (60 seconds) to allow MAC addresses learned over EVPN MPLS to remain in the MAC table before timing out. Configure the global VTEP IP address (8.8.8.8) which serves as the global identifier for MPLS encapsulation and decapsulation within the network, facilitating proper communication and tunnel establishment.

```
(config)#evpn mpls mac-ageing-time 60
(config)#evpn mpls vtep-ip-global 8.8.8.8
```

3. Define MPLS identifier (203) to support hierarchical connectivity and traffic control within the EVPN MPLS network. On the EVPN MPLS node, specify EVPN-BGP as the host reachability protocol for the specified VRF (vrf103) to communicate and exchange reachability information within the network. To enable EVPN E-Tree SC-1 on Leaf nodes, configure etree-leaf along with the MPLS identifier. This allows for efficient replication of traffic at the ingress point, optimizing the functionality of E-Tree Leaf nodes within the network architecture.

```
(config)#evpn mpls id 203 etree-leaf
(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf103
(config-evpn-mpls)#exit
```

4. Enable port-VLAN mapping (po1) with VLAN ID (103) to facilitate multi-homed access. Enable EVPN functionality on the interface, allowing it to participate in MAC address distribution across the network.

```
(config)#interface po1.103 switchport
(config-if)#encapsulation dot1q 103
(config-if)#load-interval 30
(config-access-if)#access-if-evpn
(config-access-if)#exit
```

Validation

Use the show commands described in this section to verify the network for proper MPLS EVPN E-Tree SC-1 configuration.

Verify LDP sessions on all leaf and root nodes by using the `show ldp session` command. The `state` field (OPERATIONAL) indicates that the LDP session between the device and its peers is currently active.

```
LEAF1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 24.24.24.24 | xe11 | Passive | OPERATIONAL | 30 | 01:13:29 |
| | 9.9.9.9 | xe20 | Passive | OPERATIONAL | 30 | 01:13:29 |

Verify RSVP sessions on all leaf and root nodes by using the `show rsvp session` command. The `State` field (UP) indicates that the RSVP session between the ingress and egress routers is active and operational. Identify the different paths established within the network using the `LSPName` field.

```
LEAF1#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
```

* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|---------|----------|---------|----------|--------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 9.9.9.9 | 8.8.8.8 | 5001 | 2201 | PRI | LEAF1-ROOT1- | |
| Primary | UP | 01:13:16 | 1 1 SE | - | 25601 | |
| 24.24.24.24 | 8.8.8.8 | 5002 | 2202 | PRI | LEAF1-ROOT2- | |
| Primary | UP | 01:13:05 | 1 1 SE | - | 25601 | |

Total 2 displayed, Up 2, Down 0.

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|---------|-------------|----------|---------|----------|--------------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 8.8.8.8 | 9.9.9.9 | 5001 | 2201 | PRI | ROOT1-LEAF1- | |
| Primary | UP | 01:13:45 | 1 1 SE | 25600 | - | |
| 8.8.8.8 | 24.24.24.24 | 5001 | 2201 | PRI | ROOT2-LEAF1- | |
| Primary | UP | 01:13:24 | 1 1 SE | 25601 | - | |

Total 2 displayed, Up 2, Down 0.

Verify the BGP session status on all leaf and root nodes, using the `show bgp l2vpn evpn summary` command output. The Up/Down field indicates the duration for which the BGP session has been up or down.

```
LEAF1#show bgp l2vpn evpn summary
BGP router identifier 8.8.8.8, local AS number 65535
BGP table version is 33
1 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | AD | MACIP | MCAST | ESI | PREFI |
|-------------|---|-------|--------|--------|--------|-----|------|----------|--------------|----|-------|-------|-----|-------|
| X-ROUTE | | | | | | | | | | | | | | |
| 9.9.9.9 | 4 | 65535 | 514 | 443 | 33 | 0 | 0 | 01:13:53 | 114 | 59 | 5 | 50 | 0 | |
| 0 | | | | | | | | | | | | | | |
| 24.24.24.24 | 4 | 65535 | 504 | 443 | 33 | 0 | 0 | 01:13:54 | 109 | 59 | 0 | 50 | 0 | |
| 0 | | | | | | | | | | | | | | |
| 26.26.26.26 | 4 | 65535 | 322 | 391 | 33 | 0 | 0 | 01:13:23 | 49 | 0 | 0 | 49 | 0 | |
| 0 | | | | | | | | | | | | | | |
| 29.29.29.29 | 4 | 65535 | 197 | 392 | 33 | 0 | 0 | 01:13:54 | 6 | 0 | 0 | 6 | 0 | |
| 0 | | | | | | | | | | | | | | |

Total number of neighbors 4

Total number of Established sessions 4

Verify ESI information and the forwarding tunnel status on all leaf and root nodes, by examining the `show evpn mpls` command output. The DF- Status field displays the forwarding status as either a Designated Forwarder (DF) or Non-Designated Forwarder (Non-DF), and the ESI field displays the Ethernet Segment Identifier associated with each entry.

```
LEAF1#show evpn mpls
EVPN-MPLS Information
=====
```

Codes: NW - Network Port
AC - Access Port
(u) - Untagged

| VPN-ID | EVI-Name | EVI-Type | Type | Interface | ESI | VLAN | DF-Status | Src- |
|--------|----------|----------|------|-----------|-----|------|-----------|------|
| Addr | Dst-Addr | | | | | | | |

```

203      ----      L2      NW      ----      ----      ----      ----
      8.8.8.8      29.29.29.29
203      ----      L2      NW      ----      ----      ----      ----
      8.8.8.8      9.9.9.9
203      ----      L2      NW      ----      ----      ----      ----
      8.8.8.8      24.24.24.24
203      ----      L2      NW      ----      ----      ----      ----
      8.8.8.8      26.26.26.26
203      ----      --      AC      po1.103      00:00:00:43:21:12:34:00:00:00 ----      DF      ----
      ----
203      ----      --      AC      po2.103      00:00:00:33:33:44:44:00:00:00 ----      DF      ----
      ----

```

Total number of entries are 252

Static MAC-IP Advertisement

Configure static MAC-IP advertisement through SH and MH from Root and Leaf nodes. Advertise static MAC addresses for both IPv4 and IPv6 from all MH and SH nodes. Ensure that nodes within the same MH have identical MAC addresses configured under the port-channel access port.

Configure MH Nodes

Configure static MAC addresses for IPv4 (30.30.30.3) and IPv6 (3000::1) under the MH access-port (po1) with VLAN ID (103). Repeat the same configurations for other MH nodes using different static MAC addresses for both IPv4 and IPv6.

```

!
interface po1.103 switchport
access-if-evpn
map vpn-id 203
mac 0000.7777.9999
mac 0000.7777.6666 ip 30.30.30.3
mac 0000.7777.6666 ipv6 3001::1
!

```

Configure SH Nodes

Configure static MAC addresses for IPv4 (40.40.40.4) and IPv6 (4000::1) under the SH access-port (xe27) with VLAN ID (103). This setup ensures that SH advertises these static MAC addresses over the specified access-port. Repeat the same configurations for other SH nodes using different static MAC addresses for both IPv4 and IPv6.

```

!
interface xe27.103 switchport
encapsulation dot1q 100
load-interval 30
access-if-evpn
map vpn-id 203
mac 0000.0000.0011
mac 0000.5544.4455 ip 40.40.40.4
mac 0000.5544.4455 ipv6 4000::1
!

```

Validation

Verify the MAC table entries on MH nodes (MH1, MH2 and MH3) and the SH nodes (SH1, SH2, SH3, and SH4). MH nodes advertise their MAC addresses using the ESI values. Additionally, verify the IP addresses associated with SH

nodes for MAC advertisement.

In the `show evpn mpls mac-table` command output, the MAC entries originated from Leaf Nodes will have the `LeafFlag` field status `set`.



Notes:

- MAC IPv4 or IPv6 configured under SH Leaf node access port will be advertised to the Root nodes and other Leaf nodes.
- MAC IPv4 or IPv6 configured under an MH Leaf node access port must be symmetric and will be advertised to both the Root nodes and other leaf nodes.
- MAC IPv4 or IPv6 configured under either SH or MH Root node will be advertised to both the Root nodes and the Leaf nodes.
- The Leaf-to-Leaf communication will display MAC status and tunnel status per VNI as Leaf type. The MAC will be in the discard state in the BCM shell.

```
LEAF1#show evpn mpls mac-table
```

```
=====
                                     EVPN MPLS MAC Entries
=====
VNID Interface  VlanId  In-VlanId Mac-Addr      VTEP-IP/ESI              Type Status MAC move
AccessPortDesc LeafFlag
-----
203 pol.103  ----   ----      0000.7777.9999 00:00:00:43:21:12:34:00:00:00 Static Local  --
---- 0  ----- set
203 pol.103  ----   ----      0000.7777.6666 00:00:00:43:21:12:34:00:00:00 Static Local  ---
---- 0  ----- set

Total number of entries are : 8
```

```
ROOT1#show evpn mpls mac-table
```

```
=====
                                     EVPN MPLS MAC Entries
=====
VNID Interface  VlanId  In-VlanId Mac-Addr      VTEP-IP/ESI              Type Status  MAC move
AccessPortDesc LeafFlag
-----
203 ----      ----   ----      0000.7777.9999 00:00:00:43:21:12:34:00:00:00 Static Remote  -----
    0  ----- set
203 ----      ----   ----      0000.7777.6666 00:00:00:43:21:12:34:00:00:00 Static Remote  -----
    0  ----- set

Total number of entries are : 8
```

Use the `show evpn mpls arp-cache` command to verify the Address Resolution Protocol (ARP) cache information on all nodes. This command displays entries that map IPv4 addresses to MAC addresses within the specified EVPN ID network.

```
LEAF1#show evpn mpls arp-cache
MPLS-EVPN ARP-CACHE Information
=====
```

```
EVPN-ID  Ip-Addr      Mac-Addr      Type      Age-Out  Retries-Left
-----
203      30.30.30.3   0000.7777.6666 Static Local  ----
```


Total number of entries are 5

```
ROOT1#show evpn mpls arp-cache
MPLS-EVPN ARP-CACHE Information
=====
```

```
ARP Timeout : 570 sec Random-Jitter-Max : 200
```

| EVPN-ID | Ip-Addr | Mac-Addr | Type | Age-Out | Retries-Left |
|---------|------------|----------------|---------------|---------|--------------|
| 203 | 30.30.30.3 | 0000.7777.6666 | Static Remote | ---- | |

Total number of entries are 5

Use the `show evpn mpls nd-cache` command to verify the Neighbor Discovery (ND) cache information on all nodes. This command displays entries that map IPv6 addresses to MAC addresses within the specified EVPN ID network.

```
LEAF1#show evpn mpls nd-cache
MPLS-EVPN ND-CACHE Information
=====
```

| EVPN-ID | Ip-Addr | Mac-Addr | Type | Age-Out | Retries-Left |
|---------|---------|----------------|--------------|---------|--------------|
| 203 | 3001::1 | 0000.7777.6666 | Static Local | ---- | |

Total number of entries are 4

```
ROOT1#show evpn mpls nd-cache
MPLS-EVPN ND-CACHE Information
=====
```

| EVPN-ID | Ip-Addr | Mac-Addr | Type | Age-Out | Retries-Left |
|---------|---------|----------------|---------------|---------|--------------|
| 203 | 3001::1 | 0000.7777.6666 | Static Remote | ---- | |

Total number of entries are 4

Network Topology Running Configurations

Here are the running configurations for all nodes in the given network [Figure 57](#).

LEAF1

```
!
hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
!
evpn esi hold-time 90
!
evpn etree enable
!
evpn mpls multihoming enable
!
mac vrf vrf103
  rd 8.8.8.8:103
  route-target both 65535:103
!
evpn mpls vtep-ip-global 8.8.8.8
!
evpn mpls mac-ageing-time 60
!
evpn mpls id 203 etree-leaf
  host-reachability-protocol evpn-bgp vrf103
```

```
!  
qos enable  
!  
router ldp  
  router-id 8.8.8.8  
  fast-reroute  
  graceful-restart full  
  graceful-restart timers neighbor-liveness 120  
  graceful-restart timers max-recovery 120  
  session-protection duration 10  
  targeted-peer ipv4 9.9.9.9  
    exit-targeted-peer-mode  
  targeted-peer ipv4 24.24.24.24  
    exit-targeted-peer-mode  
  transport-address ipv4 8.8.8.8  
!  
router rsvp  
!  
interface pol  
  switchport  
  load-interval 30  
  evpn multi-homed system-mac 0000.4321.1234  
!  
interface pol.103 switchport  
  encapsulation dot1q 103  
  load-interval 30  
  access-if-evpn  
    map vpn-id 203  
    mac 0000.7777.9999  
    mac 0000.7777.6666 ip 30.30.30.3  
    mac 0000.7777.6666 ipv6 3001::1  
!  
interface lo  
  ip address 8.8.8.8/32 secondary  
  ip router isis ISIS-IGP  
!  
interface xe8  
  switchport  
!  
interface xell  
  description connected to ROOT2 int xe9  
  ip address 24.1.4.25/24  
  label-switching  
  ip router isis ISIS-IGP  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface xe20  
  description connected to ROOT1 int xe20  
  ip address 26.1.2.27/24  
  label-switching  
  ip router isis ISIS-IGP  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface xe26  
  channel-group 1 mode active  
!  
interface xe27  
  speed 10g  
!  
interface xe27.100 switchport  
  encapsulation dot1q 100  
  load-interval 30  
  access-if-evpn  
    map vpn-id 200  
    mac 0000.0000.0011  
    mac 0000.5544.4455 ip 40.40.40.4
```

```
    mac 0000.5544.4455 ipv6 4000::1
!
exit
!
router isis ISIS-IGP
 is-type level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 spf-interval-exp level-1 50 2000
 metric-style wide
 mpls traffic-eng router-id 8.8.8.8
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 10000
 fast-reroute per-prefix level-1 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
 bfd all-interfaces
 net 49.0001.0000.0000.0008.00
!
router bgp 65535
 neighbor 9.9.9.9 remote-as 65535
 neighbor 24.24.24.24 remote-as 65535
 neighbor 26.26.26.26 remote-as 65535
 neighbor 29.29.29.29 remote-as 65535
 neighbor 9.9.9.9 update-source lo
 neighbor 9.9.9.9 fall-over bfd
 neighbor 24.24.24.24 update-source lo
 neighbor 24.24.24.24 fall-over bfd
 neighbor 26.26.26.26 update-source lo
 neighbor 26.26.26.26 fall-over bfd
 neighbor 29.29.29.29 update-source lo
 neighbor 29.29.29.29 fall-over bfd
!
 address-family l2vpn evpn
  neighbor 9.9.9.9 activate
  neighbor 24.24.24.24 activate
  neighbor 26.26.26.26 activate
  neighbor 29.29.29.29 activate
 exit-address-family
!
exit
!
rsvp-path LEAF1-ROOT2 mpls
 24.1.4.24 strict
!
rsvp-path LEAF1-ROOT1 mpls
 26.1.2.26 strict
!
rsvp-trunk LEAF1-ROOT1 ipv4
 primary fast-reroute protection facility
 primary path LEAF1-ROOT1
 to 9.9.9.9
!
rsvp-trunk LEAF1-ROOT2 ipv4
 primary fast-reroute protection facility
 primary path LEAF1-ROOT2
 to 24.24.24.24
!
```

LEAF2

```
!
hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
```

```
!  
evpn esi hold-time 90  
!  
evpn mpls multihoming enable  
!  
mac vrf vrf103  
  rd 26.26.26.26:103  
  route-target both 65535:103  
!  
evpn mpls vtep-ip-global 26.26.26.26  
!  
evpn mpls mac-ageing-time 60  
!  
evpn mpls id 203 etree-leaf  
  host-reachability-protocol evpn-bgp vrf103  
!  
qos enable  
!  
router ldp  
  router-id 26.26.26.26  
  fast-reroute  
  graceful-restart full  
  graceful-restart timers neighbor-liveness 120  
  graceful-restart timers max-recovery 120  
  session-protection duration 10  
  targeted-peer ipv4 9.9.9.9  
    exit-targeted-peer-mode  
  targeted-peer ipv4 24.24.24.24  
    exit-targeted-peer-mode  
  transport-address ipv4 26.26.26.26  
!  
router rsvp  
!  
interface pol  
  switchport  
  load-interval 30  
  evpn multi-homed system-mac 0000.4321.1234  
!  
interface pol.103 switchport  
  encapsulation dot1q 103  
  load-interval 30  
  access-if-evpn  
  map vpn-id 203  
!  
interface lo  
  ip address 26.26.26.26/32 secondary  
  ip router isis ISIS-IGP  
!  
interface xe3  
  channel-group 1 mode active  
!  
interface xe5  
  description connected to ROOT1 int xe30  
  ip address 26.1.7.27/24  
  label-switching  
  ip router isis ISIS-IGP  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface xe6  
  description connected to ROOT2 int xe6  
  ip address 24.1.5.25/24  
  label-switching  
  ip router isis ISIS-IGP  
  enable-ldp ipv4  
  enable-rsvp  
!  
exit
```

```
!  
router isis ISIS-IGP  
  is-type level-1  
  ignore-lsp-errors  
  lsp-gen-interval 5  
  spf-interval-exp level-1 50 2000  
  metric-style wide  
  mpls traffic-eng router-id 26.26.26.26  
  mpls traffic-eng level-1  
  capability cspf  
  dynamic-hostname  
  fast-reroute terminate-hold-on interval 10000  
  fast-reroute per-prefix level-1 proto ipv4 all  
  fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp  
  bfd all-interfaces  
  net 49.0001.0000.0000.0026.00  
!  
router bgp 65535  
  neighbor 8.8.8.8 remote-as 65535  
  neighbor 9.9.9.9 remote-as 65535  
  neighbor 24.24.24.24 remote-as 65535  
  neighbor 29.29.29.29 remote-as 65535  
  neighbor 8.8.8.8 update-source lo  
  neighbor 8.8.8.8 fall-over bfd  
  neighbor 9.9.9.9 update-source lo  
  neighbor 9.9.9.9 fall-over bfd  
  neighbor 24.24.24.24 update-source lo  
  neighbor 24.24.24.24 fall-over bfd  
  neighbor 29.29.29.29 update-source lo  
  neighbor 29.29.29.29 fall-over bfd  
!  
  address-family l2vpn evpn  
  neighbor 8.8.8.8 activate  
  neighbor 9.9.9.9 activate  
  neighbor 24.24.24.24 activate  
  neighbor 29.29.29.29 activate  
  exit-address-family  
!  
  exit  
!  
  rsvp-path LEAF2-ROOT2 mpls  
    24.1.5.24 strict  
!  
  rsvp-path LEAF2-ROOT1 mpls  
    26.1.7.26 strict  
!  
  rsvp-trunk LEAF2-ROOT1 ipv4  
    primary fast-reroute protection facility  
    primary path LEAF2-ROOT1  
    to 9.9.9.9  
!  
  rsvp-trunk LEAF2-ROOT2 ipv4  
    primary fast-reroute protection facility  
    primary path LEAF2-ROOT2  
    to 24.24.24.24  
!  
!
```

LEAF3

```
!  
evpn mpls enable  
!  
mac vrf vrf103  
  rd 29.29.29.29:103  
  route-target both 65535:103  
!
```

```
evpn mpls vtep-ip-global 29.29.29.29
!
evpn mpls mac-ageing-time 60
!
evpn mpls id 203 etree-leaf
  host-reachability-protocol evpn-bgp vrf103
!
qos enable
!
router ldp
  router-id 29.29.29.29
  fast-reroute
  graceful-restart full
  graceful-restart timers neighbor-liveness 120
  graceful-restart timers max-recovery 120
  session-protection duration 10
  targeted-peer ipv4 9.9.9.9
    exit-targeted-peer-mode
  targeted-peer ipv4 24.24.24.24
    exit-targeted-peer-mode
  transport-address ipv4 29.29.29.29
!
router rsvp
!
interface ge0
  static-channel-group 3
!
interface ge7
  description connected to ROOT1 int xe7
  ip address 26.1.6.27/24
  label-switching
  ip router isis ISIS-IGP
  enable-ldp ipv4
  enable-rsvp
!
interface lo
  ip address 29.29.29.29/32 secondary
  ip router isis ISIS-IGP
!
interface xel2
  switchport
!
interface xel2.103 switchport
  encapsulation dot1q 103
  load-interval 30
  access-if-evpn
  map vpn-id 203
!
interface xel4
  description connected to ROOT2 int xe13
  ip address 24.1.8.25/24
  label-switching
  ip router isis ISIS-IGP
  enable-ldp ipv4
  enable-rsvp
!
  exit
!
router isis ISIS-IGP
  is-type level-1
  ignore-lsp-errors
  lsp-gen-interval 5
  spf-interval-exp level-1 50 2000
  metric-style wide
  mpls traffic-eng router-id 29.29.29.29
  mpls traffic-eng level-1
  capability cspf
  dynamic-hostname
```

```
fast-reroute terminate-hold-on interval 10000
fast-reroute per-prefix level-1 proto ipv4 all
fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
bfd all-interfaces
net 49.0001.0000.0000.0029.00
!
router bgp 65535
neighbor 8.8.8.8 remote-as 65535
neighbor 9.9.9.9 remote-as 65535
neighbor 24.24.24.24 remote-as 65535
neighbor 26.26.26.26 remote-as 65535
neighbor 8.8.8.8 update-source lo
neighbor 8.8.8.8 fall-over bfd
neighbor 9.9.9.9 update-source lo
neighbor 9.9.9.9 fall-over bfd
neighbor 24.24.24.24 update-source lo
neighbor 24.24.24.24 fall-over bfd
neighbor 26.26.26.26 update-source lo
neighbor 26.26.26.26 fall-over bfd
!
address-family l2vpn evpn
neighbor 8.8.8.8 activate
neighbor 9.9.9.9 activate
neighbor 24.24.24.24 activate
neighbor 26.26.26.26 activate
exit-address-family
!
exit
!
rsvp-path LEAF3-ROOT2 mpls
24.1.8.24 strict
!
rsvp-path LEAF3-ROOT1 mpls
26.1.6.26 strict
!
rsvp-trunk LEAF3-ROOT1 ipv4
primary fast-reroute protection facility
primary path LEAF3-ROOT1
to 9.9.9.9
!
rsvp-trunk LEAF3-ROOT2 ipv4
primary fast-reroute protection facility
primary path LEAF3-ROOT2
to 24.24.24.24
!
```

ROOT1

```
!
hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
!
evpn esi hold-time 90
!
evpn mpls multihoming enable
!
mac vrf vrf103
rd 9.9.9.9:103
route-target both 65535:103
!
evpn mpls vtep-ip-global 9.9.9.9
!
evpn mpls mac-ageing-time 60
!
evpn mpls id 203
```

```
host-reachability-protocol evpn-bgp vrf103
!
qos enable
!
bridge 1 protocol rstp vlan-bridge
!
router ldp
router-id 9.9.9.9
fast-reroute
graceful-restart full
graceful-restart timers neighbor-liveness 120
graceful-restart timers max-recovery 120
session-protection duration 10
targeted-peer ipv4 8.8.8.8
  exit-targeted-peer-mode
targeted-peer ipv4 26.26.26.26
  exit-targeted-peer-mode
transport-address ipv4 9.9.9.9
!
router rsvp
!
interface pol100
switchport
load-interval 30
evpn multi-homed system-mac 0000.1111.2222
!
interface pol100.103 switchport
encapsulation dot1q 103
load-interval 30
access-if-evpn
  map vpn-id 203
!
interface lo
ip address 9.9.9.9/32 secondary
ip router isis ISIS-IGP
!
interface xe7
description connected to LEAF3 int ge7
speed 1g
ip address 26.1.6.26/24
label-switching
ip router isis ISIS-IGP
enable-ldp ipv4
enable-rsvp
!
interface xe10
channel-group 100 mode active
!
interface xe17.100 switchport
description for Static mac advertize
encapsulation dot1q 100
load-interval 30
access-if-evpn
  map vpn-id 200
  mac 0000.0000.0022
  mac 0000.00dc.0001 ip 10.10.10.1
  mac 0000.00dc.0001 ipv6 1001::1
!
interface xe19
switchport
!
interface xe20
description connected to LEAF1 int xe20
ip address 26.1.2.26/24
label-switching
ip router isis ISIS-IGP
enable-ldp ipv4
enable-rsvp
```



```
!  
interface xe30  
  description connected to LEAF2 int xe5  
  speed 10g  
  ip address 26.1.7.26/24  
  label-switching  
  ip router isis ISIS-IGP  
  enable-ldp ipv4  
  enable-rsvp  
!  
  exit  
!  
router isis ISIS-IGP  
  is-type level-1  
  ignore-lsp-errors  
  lsp-gen-interval 5  
  spf-interval-exp level-1 50 2000  
  metric-style wide  
  mpls traffic-eng router-id 9.9.9.9  
  mpls traffic-eng level-1  
  capability cspf  
  dynamic-hostname  
  fast-reroute terminate-hold-on interval 10000  
  fast-reroute per-prefix level-1 proto ipv4 all  
  fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp  
  bfd all-interfaces  
  net 49.0001.0000.0000.0009.00  
!  
router bgp 65535  
  neighbor 8.8.8.8 remote-as 65535  
  neighbor 24.24.24.24 remote-as 65535  
  neighbor 26.26.26.26 remote-as 65535  
  neighbor 29.29.29.29 remote-as 65535  
  neighbor 8.8.8.8 update-source lo  
  neighbor 8.8.8.8 fall-over bfd  
  neighbor 24.24.24.24 update-source lo  
  neighbor 24.24.24.24 fall-over bfd  
  neighbor 26.26.26.26 update-source lo  
  neighbor 26.26.26.26 fall-over bfd  
  neighbor 29.29.29.29 update-source lo  
  neighbor 29.29.29.29 fall-over bfd  
  !  
  address-family l2vpn evpn  
  neighbor 8.8.8.8 activate  
  neighbor 24.24.24.24 activate  
  neighbor 26.26.26.26 activate  
  neighbor 29.29.29.29 activate  
  exit-address-family  
  !  
  exit  
!  
  rsvp-path ROOT1-LEAF3 mpls  
    26.1.6.27 strict  
  !  
  rsvp-path ROOT1-LEAF2 mpls  
    26.1.7.27 strict  
  !  
  rsvp-path ROOT1-LEAF1 mpls  
    26.1.2.27 strict  
  !  
  rsvp-trunk ROOT1-LEAF1 ipv4  
    primary fast-reroute protection facility  
    primary path ROOT1-LEAF1  
    to 8.8.8.8  
  !  
  rsvp-trunk ROOT1-LEAF2 ipv4  
    primary fast-reroute protection facility  
    primary path ROOT1-LEAF2
```

```
to 26.26.26.26
!
rsvp-trunk ROOT1-LEAF3 ipv4
primary fast-reroute protection facility
primary path ROOT1-LEAF3
to 29.29.29.29
!
```

ROOT2

```
!
hardware-profile filter evpn-mpls-mh enable
!
evpn mpls enable
!
evpn esi hold-time 90
!
evpn mpls multihoming enable
!
mac vrf vrf103
rd 24.24.24.24:103
route-target both 65535:103
!
evpn mpls vtep-ip-global 24.24.24.24
!
evpn mpls mac-ageing-time 60
!
evpn mpls id 203
host-reachability-protocol evpn-bgp vrf103
!
qos enable
!
router ldp
router-id 24.24.24.24
fast-reroute
graceful-restart full
graceful-restart timers neighbor-liveness 120
graceful-restart timers max-recovery 120
session-protection duration 10
targeted-peer ipv4 8.8.8.8
exit-targeted-peer-mode
targeted-peer ipv4 26.26.26.26
exit-targeted-peer-mode
transport-address ipv4 24.24.24.24
!
router rsvp
!
interface pol00
switchport
load-interval 30
evpn multi-homed system-mac 0000.1111.2222
!
interface pol00.103 switchport
encapsulation dot1q 103
load-interval 30
access-if-evpn
map vpn-id 203
!
interface lo
ip address 24.24.24.24/32 secondary
ip router isis ISIS-IGP
!
interface xe6
description connected to LEAF2 int xe6
speed 10g
ip address 24.1.5.24/24
```

```
label-switching
ip router isis ISIS-IGP
enable-ldp ipv4
enable-rsvp
!
interface xe7
  switchport
!
interface xe9
  description connected to LEAF1 int xe11
  speed 10g
  ip address 24.1.4.24/24
  label-switching
  ip router isis ISIS-IGP
  enable-ldp ipv4
  enable-rsvp
!
interface xe13
  description connected to LEAF3 int xe14
  speed 10g
  ip address 24.1.8.24/24
  label-switching
  ip router isis ISIS-IGP
  enable-ldp ipv4
  enable-rsvp
!
interface xe18
  channel-group 100 mode active
!
exit
!
router isis ISIS-IGP
  is-type level-1
  ignore-lsp-errors
  lsp-gen-interval 5
  spf-interval-exp level-1 50 2000
  metric-style wide
  mpls traffic-eng router-id 24.24.24.24
  mpls traffic-eng level-1
  capability cspf
  dynamic-hostname
  fast-reroute terminate-hold-on interval 10000
  fast-reroute per-prefix level-1 proto ipv4 all
  fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
  bfd all-interfaces
  net 49.0001.0000.0000.0024.00
!
router bgp 65535
  neighbor 8.8.8.8 remote-as 65535
  neighbor 9.9.9.9 remote-as 65535
  neighbor 26.26.26.26 remote-as 65535
  neighbor 29.29.29.29 remote-as 65535
  neighbor 8.8.8.8 update-source lo
  neighbor 8.8.8.8 fall-over bfd
  neighbor 9.9.9.9 update-source lo
  neighbor 9.9.9.9 fall-over bfd
  neighbor 26.26.26.26 update-source lo
  neighbor 26.26.26.26 fall-over bfd
  neighbor 29.29.29.29 update-source lo
  neighbor 29.29.29.29 fall-over bfd
!
  address-family l2vpn evpn
  neighbor 8.8.8.8 activate
  neighbor 9.9.9.9 activate
  neighbor 26.26.26.26 activate
  neighbor 29.29.29.29 activate
  exit-address-family
!
```

```
exit
!
rsvp-path ROOT2-LEAF1 mpls
 24.1.4.25 strict
!
rsvp-path ROOT2-LEAF2 mpls
 24.1.5.25 strict
!
rsvp-path ROOT2-LEAF3 mpls
 24.1.8.25 strict
!
rsvp-trunk ROOT2-LEAF1 ipv4
 primary fast-reroute protection facility
 primary path ROOT2-LEAF1
 to 8.8.8.8
!
rsvp-trunk ROOT2-LEAF2 ipv4
 primary fast-reroute protection facility
 primary path ROOT2-LEAF2
 to 26.26.26.26
!
rsvp-trunk ROOT2-LEAF3 ipv4
 primary fast-reroute protection facility
 primary path ROOT2-LEAF3
 to 29.29.29.29
!
```

CE SWITCH

```
!
bridge 1 protocol rstp vlan-bridge
!
vlan database
 vlan-reservation 4030-4094
 vlan 2-3010 bridge 1 state enable
!
interface po100
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
!
interface lo
 ip address 32.32.32.32/32 secondary
!
interface xe9
 channel-group 100 mode active
!
interface xe17
 channel-group 100 mode active
!
interface xe1
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
!
exit
!
```

SWITCH1

```
!
bridge 1 protocol rstp vlan-bridge
!
vlan-reservation 4020-4062
```

```

vlan 2-3000 bridge 1 state enable
!
interface po1
switchport
bridge-group 1
switchport mode trunk
switchport trunk allowed vlan all
!
interface lo
ip address 7.7.7.7/32 secondary
!
interface xe1/1
switchport
bridge-group 1
switchport mode trunk
switchport trunk allowed vlan all
!
interface xe3/1
channel-group 1 mode active
!
interface xe3/3
channel-group 1 mode active
!
exit
!
```

SWITCH2

```

!
bridge 1 protocol rstp vlan-bridge
!
vlan database
vlan 2-3000 bridge 1 state enable
!
interface sa3
switchport
bridge-group 1
switchport mode trunk
switchport trunk allowed vlan all
!
interface ge0
static-channel-group 3
!
interface lo
ip address 23.23.23.23/32 secondary
!
interface xe4
switchport
bridge-group 1
switchport mode trunk
switchport trunk allowed vlan all
!
```

E-Tree Active-Standby Configuration

To set up an E-Tree Scenario 1 network with Active-Standby redundancy and load balancing, follow these steps:

- Connect the Switch (P1) to the Root1, LEAF1, and LEAF2 nodes in the [Figure 57](#).
- Set up the VRF, EVPN, Port-Active, and Single-Active Redundancy configuration on Root MH and Leaf MH nodes.

For more details on Active-Standby configuration, refer to the section [EVPN Active-Standby \(page 1752\)](#).

LEAF1

```
!  
mac vrf vrf600  
  rd 26.26.26.26:600  
  route-target both 65535:600  
!  
evpn mpls id 681 etree-leaf  
  host-reachability-protocol evpn-bgp vrf600  
!  
interface pol  
  switchport  
  load-interval 30  
  evpn multi-homed system-mac 0000.4321.1234 load-balancing port-active  
  service-carving auto  
!  
interface pol.681 switchport  
  encapsulation dot1q 681  
  load-interval 30  
  access-if-evpn  
  map vpn-id 681  
!  
interface sal  
  switchport  
  load-interval 30  
  evpn multi-homed esi 11:22:33:00:00:00:55:66:77 load-balancing single-active  
  service-carving auto  
!  
interface sal.681 switchport  
  encapsulation dot1q 681  
  load-interval 30  
  access-if-evpn  
  map vpn-id 681  
!  
interface xe4  
  description connected to P1 int xe43  
  speed 10g  
  load-interval 30  
  ip address 25.1.2.25/24  
  label-switching  
  ip router isis ISIS-IGP  
  enable-ldp ipv4  
  enable-rsvp  
!
```

LEAF2

```
!  
mac vrf vrf600  
  rd 26.26.26.26:600  
  route-target both 65535:600  
!  
evpn mpls id 681 etree-leaf  
  host-reachability-protocol evpn-bgp vrf600  
!  
interface pol  
  switchport  
  load-interval 30  
  evpn multi-homed system-mac 0000.4321.1234 load-balancing port-active  
  service-carving auto  
!  
interface pol.681 switchport  
  encapsulation dot1q 681  
  load-interval 30  
  access-if-evpn  
  map vpn-id 681
```

```

!
interface sa2
  switchport
  load-interval 30
  evpn multi-homed esi 11:22:33:00:00:00:55:66:77 load-balancing single-active
  service-carving auto
!
interface sa2.681 switchport
  encapsulation dot1q 681
  load-interval 30
  access-if-evpn
  map vpn-id 681
!
interface xe21
  description connected to P1 int xe43
  speed 10g
  load-interval 30
  ip address 27.1.2.25/24
  label-switching
  ip router isis ISIS-IGP
  enable-ldp ipv4
  enable-rsvp
!

```

P1

```

!
router ldp
  router-id 6.6.6.6
  graceful-restart full
  graceful-restart timers neighbor-liveness 120
  graceful-restart timers max-recovery 120
  session-protection duration 10
  transport-address ipv4 6.6.6.6
!
interface lo
  ip address 127.0.0.1/8
  ip address 6.6.6.6/32 secondary
  ipv6 address ::1/128
  ip router isis ISIS-IGP
!
interface xe43
  description connected to LEAF1 int xe4
  speed 10g
  load-interval 30
  ip address 25.1.2.24/24
  label-switching
  ip router isis ISIS-IGP
  enable-ldp ipv4
  enable-rsvp
!
interface xe45
  description connected to ROOT1 int xe2
  speed 10g
  load-interval 30
  ip address 26.1.3.27/24
  label-switching
  ip router isis ISIS-IGP
  enable-ldp ipv4
  enable-rsvp
!
interface xe47
  description connected to LEAF2 int xe21
  speed 10g
  load-interval 30
  ip address 27.1.2.24/24
  label-switching
  ip router isis ISIS-IGP

```

```

enable-ldp ipv4
enable-rsvp
!
exit
!
router isis ISIS-IGP
 is-type level-1
 authentication mode md5 level-1
 ignore-lsp-errors
 lsp-gen-interval 5
 spf-interval-exp level-1 50 2000
 metric-style wide
 mpls traffic-eng router-id 6.6.6.6
 mpls traffic-eng level-1
 capability cspf
 dynamic-hostname
 fast-reroute terminate-hold-on interval 10000
 fast-reroute per-prefix level-1 proto ipv4 all
 fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
 bfd all-interfaces
 net 49.0001.0000.0000.0006.00
!
```

Validation

To verify the status of the ESI, whether it's active or standby, use the `show evpn load-balance all` command. This command helps debug and understand if the election process is occurring correctly. For the ESI 00:00:00:43:21:12:34:00:00:00, LEAF1 is active, and LEAF2 is on standby in port-active mode. For the ESI 00:11:22:33:00:00:00:55:66:77, LEAF2 is active, and LEAF1 is on standby in single-active mode.

LEAF1#show evpn load-balance all

| ESI | Revertive | AC-DF | Status | AC-IF/PE | PE-IP-ADDRESS | Redundancy | Service-carving weight |
|-------------------------------|-----------|-------|--------|----------|---------------|------------|------------------------|
| 00:00:00:43:21:12:34:00:00:00 | active | auto | 0 | LOCAL | 8.8.8.8 | port- | |
| 00:00:00:43:21:12:34:00:00:00 | active | auto | 0 | REMOTE | 26.26.26.26 | port- | |
| 00:11:22:33:00:00:00:55:66:77 | active | auto | 0 | sa1.681 | 8.8.8.8 | single- | |

LEAF2#show evpn load-balance all

| ESI | Revertive | AC-DF | Status | AC-IF/PE | PE-IP-ADDRESS | Redundancy | Service-carving weight |
|-------------------------------|-----------|-------|--------|----------|---------------|------------|------------------------|
| 00:00:00:43:21:12:34:00:00:00 | active | auto | 0 | REMOTE | 8.8.8.8 | port- | |
| 00:00:00:43:21:12:34:00:00:00 | active | auto | 0 | LOCAL | 26.26.26.26 | port- | |
| 00:11:22:33:00:00:00:55:66:77 | active | auto | 0 | sa2.681 | 26.26.26.26 | single- | |

All MAC addresses in Root and Leaf nodes will be synchronized.

LEAF1#show evpn mpls mac-table

| VNID | Interface | VlanId | In-VlanId | Mac-Addr | VTEP- |
|-----------------------|-----------|--------|-----------|----------|-------|
| EVPN MPLS MAC Entries | | | | | |


```

Ip/ESI                                     Type          Status      MAC move AccessPortDesc LeafFlag
-----
681      po1.681      ----      ----      0000.da00.0001
00:00:00:43:21:12:34:00:00:00      Dynamic Local      -----      0      -----
set
681      ----      ----      ----      0000.ea00.0001
00:00:00:11:11:22:22:00:00:00      Dynamic Remote      -----      0      -----      --
--

Total number of entries are : 2

```

```
LEAF2#show evpn mpls mac-table
```

```

=====
EVPN MPLS MAC Entries
=====
VNID      Interface VlanId      In-VlanId Mac-Addr      VTEP-
Ip/ESI                                     Type          Status      MAC move AccessPortDesc LeafFlag
-----
681      ----      ----      ----      0000.da00.0001
00:00:00:43:21:12:34:00:00:00      Dynamic Remote      -----      0      -----
set
681      ----      ----      ----      0000.ea00.0001
00:00:00:11:11:22:22:00:00:00      Dynamic Remote      -----      0      -----      --
--

Total number of entries are : 2

```

```
ROOT1#show evpn mpls mac-table
```

```

=====
EVPN MPLS MAC Entries
=====
VNID      Interface VlanId      In-VlanId Mac-Addr      VTEP-
Ip/ESI                                     Type          Status      MAC move AccessPortDesc LeafFlag
-----
681      ----      ----      ----      0000.da00.0001
00:00:00:43:21:12:34:00:00:00      Dynamic Remote      -----      0      -----
set
681      po100.681 ----      ----      0000.ea00.0001
00:00:00:11:11:22:22:00:00:00      Dynamic Local      -----      0      -----      --
--

Total number of entries are : 2

```

Implementation Examples

Here is an example scenario and a solution for implementing E-Tree Scenario 1.

Use Case 1: Specific traffic isolation and control measures are essential in a network of EVPN L2VPN services or instances. Within a broadcast domain, services communicating with each other may result in flooding BUM traffic to all services within the domain. Moreover, hosts are learned and advertised between different sites/services.

Solution 1: Implementing an EVPN E-Tree Scenario 1 solution defines the network topology with distinct Root and Leaf classifications, BUM traffic flooding can be minimized, and traffic isolation can be achieved. This ensures

efficient communication between services while preventing unnecessary traffic propagation and maintaining network integrity.

Use Case 2: An Internet Service Provider (ISP) provides services to multiple subscribers and aims to facilitate communication with them. However, the ISP needs to ensure that subscribers exclusively communicate with the ISP and not among themselves.

Solution 2: Implementing EVPN E-Tree Scenario 1 is essential to fulfill this requirement. By categorizing ISP services as Root and subscribers as Leaf, traffic isolation can be enforced. This configuration enables the ISP to communicate with subscribers while preventing inter-subscriber communication. As a result, network security is enhanced, and the ISP maintains control over communication within its network.

E-Tree Scenario 1 Commands

The EVPN E-Tree Scenario 1 introduces the [evpn etree \(page 1920\)](#) configuration command in OcNOS.

Revised Commands

evpn mpls id

Introduced `etree-leaf` parameter in the command `evpn mpls id <ID> etree-leaf`, which allows users to tailor MPLS EVPN behavior on a network device, indicating its participation as a leaf node in an E-Tree Scenario 1 deployment. For more details, refer to the [evpn mpls id \(page 1936\)](#) command in the [EVPN MPLS Commands \(page 1538\)](#) section.

Troubleshooting

1. When traffic, whether unicast (UC) or broadcast, is passed to the Intra Leaf site:
 - Check the sub-interface or physical interface counters to monitor traffic throughput and potential issues.
 - Verify the Leaf status of the corresponding VNI to ensure proper functionality.
 - Use packet sniffing tools to analyze packets in the egress direction for any anomalies or errors.
 - MAC entries learned via leaf access port should include the `set` keyword in the MAC table output.
2. If UC traffic is routed within inter-PE leaf sites:
 - Check the Leaf status of the VNI at both participating PE devices to confirm operational status.
 - Check if the advertised MAC is in discard or non-discard status using the `show mac table` command and `12 show` in the BCM shell.
3. Verify if BUM traffic is transmitted between Leaf sites inter-PE:
 - Ensure that a BUM tunnels are not established between inter-PE devices.
 - Validate this by examining the Multicast ingress group, using the `show evpn mpls tunnel` command. For EVPN MPLS, confirm that BUM tunnels are not created.
4. Investigate UC traffic drops from the Root to MH Leaf PE:
 - Check if MAC addresses are not installed in discard status within the MH peer's access port. This status could indicate issues with MAC learning or forwarding.
5. Evaluate traffic between Root and Leaf:
 - Confirm the establishment of both UC and BUM tunnels.

- Ensure that unicast MAC addresses are not marked with a discard status in the MAC table.
6. Validate the exchange of routes between two BGP L2VPN peers:
 - Monitor BGP (Border Gateway Protocol) sessions to verify successful route exchange and propagation between the peers.
 7. Convergence: Assess convergence by checking BFD configuration between BGP sessions.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms or Acronym | Description |
|--|--|
| Ethernet VPN Ethernet-Tree (EVPN E-Tree) | A networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, categorizing network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them. |
| Ethernet Virtual Private Network (EVPN) | A Layer 2 VPN technology that extends Ethernet services across data centers and wide-area networks using BGP. |
| Multi-homing (MH) | The ability of a device to connect to multiple network segments simultaneously to increase network availability and redundancy. |
| Provider Edge (PE) Node | A device at the edge of a service provider network that connects to customer premises equipment (CE) and participates in providing services to customers. |
| Leaf Node | In the context of EVPN E-Tree, a network node categorized to handle communication within specific broadcast domains and may connect to Root nodes. |
| Root Node | A network node within EVPN E-Tree that serves as the central point of communication and handles BUM traffic distribution. |
| Ethernet Segment Identifier (ESI) | A unique identifier used to identify Ethernet segments within a MPLS network. |

EVPN MPLS E-Tree Scenario 2

Overview

Ethernet VPN Ethernet-Tree (EVPN E-Tree) Scenario 2 (SC-2) [RFC 8317] enables Root and Leaf sites to co-exist on the same Provider Edge (PE) device within a single EVPN Instance (EVI). The solution enforces traffic filtering rules both within a PE (intra-PE) and across PEs (inter-PE), ensuring strict Leaf-to-Leaf communication control while maintaining flexible service deployment.

Scenario 2 supports both single-homed (SH) and multi-homed (MH) configurations and applies to Qumran2 (Q2) series platforms. All other EVPN E-Tree functionality and route types conform to [RFC 7432], except where extended for Scenario 2 behavior.

Inter-PE E-tree: New Route Advertisements

Scenario 2 introduces the following new route advertisements to support Leaf traffic isolation across PEs:

- **Leaf Host Advertisement (RT-2):** RT-2 advertisements for Leaf hosts include a new E-Tree Extended Community attribute. This attribute identifies the host as a Leaf for unicast traffic handling.
- **EAD per ES Advertisement (RT-1):** A new EAD per ES Route Type-1 (RT-1) is advertised with ESI=0. This route carries the E-Tree Extended Community attribute, which contains the Leaf Label. This mechanism is used to filter unknown traffic between a leaf source and a leaf destination.

Feature Characteristics

EVPN E-Tree Scenario 2 allows Root and Leaf Attachment Circuits (ACs) to exist within the same EVI on a PE. Traffic between Leaf ACs is restricted both locally and across PEs, maintaining E-Tree hierarchy and service separation.

Intra-PE Traffic Filtering

OcNOS filters traffic locally between Leaf ACs that belong to the same EVI on a single PE.

Traffic Flow (Unicast or BUM):

Leaf AC to Leaf AC (same PE): Traffic originating from one Leaf AC and destined for another Leaf AC on the same PE is dropped. This is due to the enforcement of split-horizon, which prevents Leaf-to-Leaf communication within the same PE. Any unicast or unknown traffic is filtered based on the source and destination ACs; if both are Leaf ACs, the traffic is dropped.

Inter-PE Traffic Filtering (Leaf ACs in the Same EVI)

To restrict communication between Leaf Access Circuits (ACs) located on different Provider Edges (PEs) within the same EVPN Instance (EVI), the amended routes described previously are utilized.

Unicast Traffic Filtering

Traffic filtering for Unicast traffic between Leaf ACs is performed at the originating Leaf AC PE node (the ingress PE node). This is possible because the host advertised from a Leaf AC at a remote PE is identified as a Leaf through the new RT-2 E-TREE attribute. Consequently, the Local PE recognizes the remote Host as one advertised from a Leaf AC, enabling ingress filtering. Reference: [RFC 8317, Section 4.1](#).

BUM Traffic Filtering

- Traffic filtering for Unknown Broadcast, Unknown Unicast, Multicast (BUM) traffic between Leaf ACs is implemented at the egress PE device. In the case of BUM traffic, the receivers include all ACs in the EVPN instance, some of which may be Root and others Leaf. Therefore, filtering cannot be restricted at the ingress PE.
- The remote PE uses RT-1 with ESI=0 and the E-TREE attribute to advertise a Leaf Label. The local Leaf AC uses this Leaf Label when sending BUM traffic.

Reference: [RFC 8317, Section 4.2](#)

Handling Multi-Homing (MH)

In Multi-Homing scenarios, if a Leaf AC is also on an ESI Multi-Homing port, the Leaf Label is prioritized over the ESI Label for advertising.

Intra-PE Traffic Flow Details

- **Co-existing Root and Leaf Sites on the Same PE:** A given Provider Edge (PE) may simultaneously host both Root and Leaf Attachment Circuits (ACs) for a specific Ethernet Virtual Instance (EVI).
- **Ingress Filtering (Unicast or Known Traffic):**
 - Traffic is subjected to filtering upon entering the PE via a Leaf AC.
 - A dedicated grouping identifies all Media Access Control (MAC) addresses learned from remote Leaf ACs.
 - Leaf-to-Leaf unicast traffic is discarded to prevent unauthorized communication.
- **Egress Filtering (BUM Traffic):**
 - BUM traffic includes a Leaf Label, which is advertised by the remote PE.
 - The Egress PE utilizes the Leaf Label to inhibit traffic transmission towards Leaf ACs.
 - This mechanism is effective for both single-homed (SH) and multi-homed (MH) Leaf ACs.
- **Route Exchange and Attributes:**

| Route Type | Attribute Function | Filtering Scope |
|-----------------------------------|---|-----------------|
| RT-2 (MAC/IP Advertisement) | Carries Leaf indication for ingress unicast filtering | Ingress PE |
| RT-1 (Ethernet A-D per ES, ESI=0) | Carries Leaf Label for egress BUM filtering | Egress PE |



Note: The Leaf Label is scoped per PE, not per EVI or per ES.

Traffic Filtering Rules for All Combinations

Unicast

- **Leaf AC to Leaf MAC:** Traffic originating from a Leaf AC and destined for a MAC address learned from any Leaf AC on a remote PE is dropped, thereby preventing Leaf-to-Leaf communication across PEs.
- **Root AC to Leaf MAC:** Traffic is permitted, facilitating Root-to-Leaf communication.
- **Root AC to non-Leaf (Root) MAC:** Traffic is permitted, facilitating Root-to-Root communication.
- **Leaf AC to non-Leaf (Root) MAC:** Traffic is permitted, facilitating Leaf-to-Root communication.

BUM

- **SH Leaf AC:** Traffic is tagged with the Leaf Label; the egress PE drops it toward Leaf ACs.
- **MH Leaf AC on ESI port:** Traffic carries the Leaf Label; the egress PE blocks delivery to all Leaf ACs (both SH and MH).
- **Leaf-to-Leaf traffic within the same PE:** Traffic is dropped (via split-horizon functionality).



Note: Configure `hardware-profile filter evpn-mpls-mh` group even for single-homing nodes to enable Leaf Label enforcement.

Benefits

- **Strict traffic enforcement:** Leaf-to-Leaf communication is always blocked, maintaining E-Tree hierarchy and isolation.
- **Flexible service deployment:** Root and Leaf sites can co-exist on the same PE, simplifying design and reducing the number of EVIs.
- **Simplified control-plane operation:** A single route type per EVI with the E-Tree Extended Community and Leaf Label reduces BGP processing complexity.
- **Enhanced traffic security:** Ingress and egress filtering ensure unauthorized or misrouted traffic is dropped at the hardware level.
- **Consistent SH or MH operation:** Unified filtering logic supports both single-homed and multi-homed Leaf ACs.
- **Efficient hardware utilization:** Leaf Label allocation per PE enables hardware-level enforcement without additional per-EVI or per-ES labels.

E-Tree Scenario 2 Prerequisites

Before configuring E-Tree Scenario-2 in [Figure 58](#) topology, ensure the following prerequisites:

Root PEs and Leaf PEs**Underlay Transport Setup (MPLS, LDP, RSVP, IGP)**

MPLS, LDP, RSVP, and IS-IS (or OSPF) are operational across all PE and core routers. These protocols establish the MPLS underlay for label distribution and traffic forwarding.

```
!
router ospf 100
  ospf router-id 1.1.1.1
  bfd all-interfaces
  network 1.1.1.1/32 area 0.0.0.0
  network 12.1.1.0/24 area 0.0.0.0
  network 13.1.1.0/24 area 0.0.0.0
!
evpn mpls vtep-ip-global 1.1.1.1
!
evpn mpls id 100
  host-reachability-protocol evpn-bgp vrf100
!
evpn mpls id 101
```

```

    host-reachability-protocol evpn-bgp vrf101
    !
    evpn mpls id 102
    host-reachability-protocol evpn-bgp vrf102
    !
    router ldp
    graceful-restart full
    targeted-peer ipv4 2.2.2.2
    exit-targeted-peer-mode
    targeted-peer ipv4 3.3.3.3
    exit-targeted-peer-mode
    targeted-peer ipv4 4.4.4.4
    exit-targeted-peer-mode
    transport-address ipv4 1.1.1.1
    !
    router rsvp
    !
    rsvp-path PE1-PE3 mpls
    13.1.1.2 strict
    17.1.1.1 strict
    19.1.1.2 strict
    !
    rsvp-path PE1-PE4 mpls
    12.1.1.2 strict
    17.1.1.2 strict
    21.1.1.2 strict
    !
    rsvp-trunk PE1-PE3 ipv4
    primary fast-reroute protection one-to-one
    primary path PE1-PE3
    to 3.3.3.3
    !
    rsvp-trunk PE1-PE4 ipv4
    primary fast-reroute protection facility
    primary path PE1-PE4
    to 4.4.4.4
    !
    rsvp-trunk PE1-PE2 ipv4
    to 2.2.2.2
    !

```

Loopback Interface Reachability

Loopback interfaces on each PE and core router are reachable through the IGP. Loopbacks provide router IDs for BGP EVPN and LDP sessions.

```

!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
 ip router isis ISIS-IGP
 enable-rsvp
!

```

Global EVPN and Multihoming Enablement

EVPN MPLS and Multihoming are globally enabled and allow multi-homed Ethernet segments for redundancy.

```

!
evpn mpls enable
!
evpn esi hold-time 200
!
evpn mpls multihoming enable
!

```

BGP EVPN Control Plane Establishment

BGP EVPN sessions are established among all PE routers. Each PE's loopback interface is used for BGP EVPN update-source and RSVP destination.

```
!  
router bgp 65010  
  neighbor EVPN peer-group  
  neighbor EVPN remote-as 65010  
  neighbor EVPN update-source lo  
  neighbor EVPN advertisement-interval 0  
  neighbor EVPN fall-over bfd multihop  
  neighbor 2.2.2.2 peer-group EVPN  
  neighbor 3.3.3.3 peer-group EVPN  
  neighbor 4.4.4.4 peer-group EVPN  
!  
address-family l2vpn evpn  
  neighbor EVPN activate  
exit-address-family  
!  
exit  
!
```

Core-Facing Interface Configuration

Port-channels or physical links toward the core have MPLS, LDP, RSVP, and IGP enabled. Each core-facing link supports transport signaling for MPLS services.

```
!  
interface po10  
  description connected-to-p1  
  load-interval 30  
  ip address 12.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface po20  
  description connected-to-p2  
  load-interval 30  
  ip address 13.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!
```

MPLS Core Fabric Formation (P1 or P2 Routers)

[Core Routers \(page 1912\)](#) (P1 and P2) form the MPLS transport fabric between the root and leaf PEs (Core routers interconnect the EVPN PEs and carry labeled traffic.).

```
!  
interface xe1  
  description connected-to-p2  
  load-interval 30  
  ip address 17.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4
```



```
enable-rsvp
!
```

Access Switch VLAN Trunk Configuration

[Access Switches \(page 1917\)](#) (SW-1 and SW-2) are configured with VLAN trunks toward PEs carrying tagged leaf traffic.

```
!
interface po1000
  switchport
  bridge-group 1 spanning-tree disable
  switchport mode trunk
  switchport trunk allowed vlan all
!
interface xe3/1
  port breakout enable 4X10g
  switchport
  bridge-group 1 spanning-tree disable
  switchport mode trunk
  switchport trunk allowed vlan all
!
```

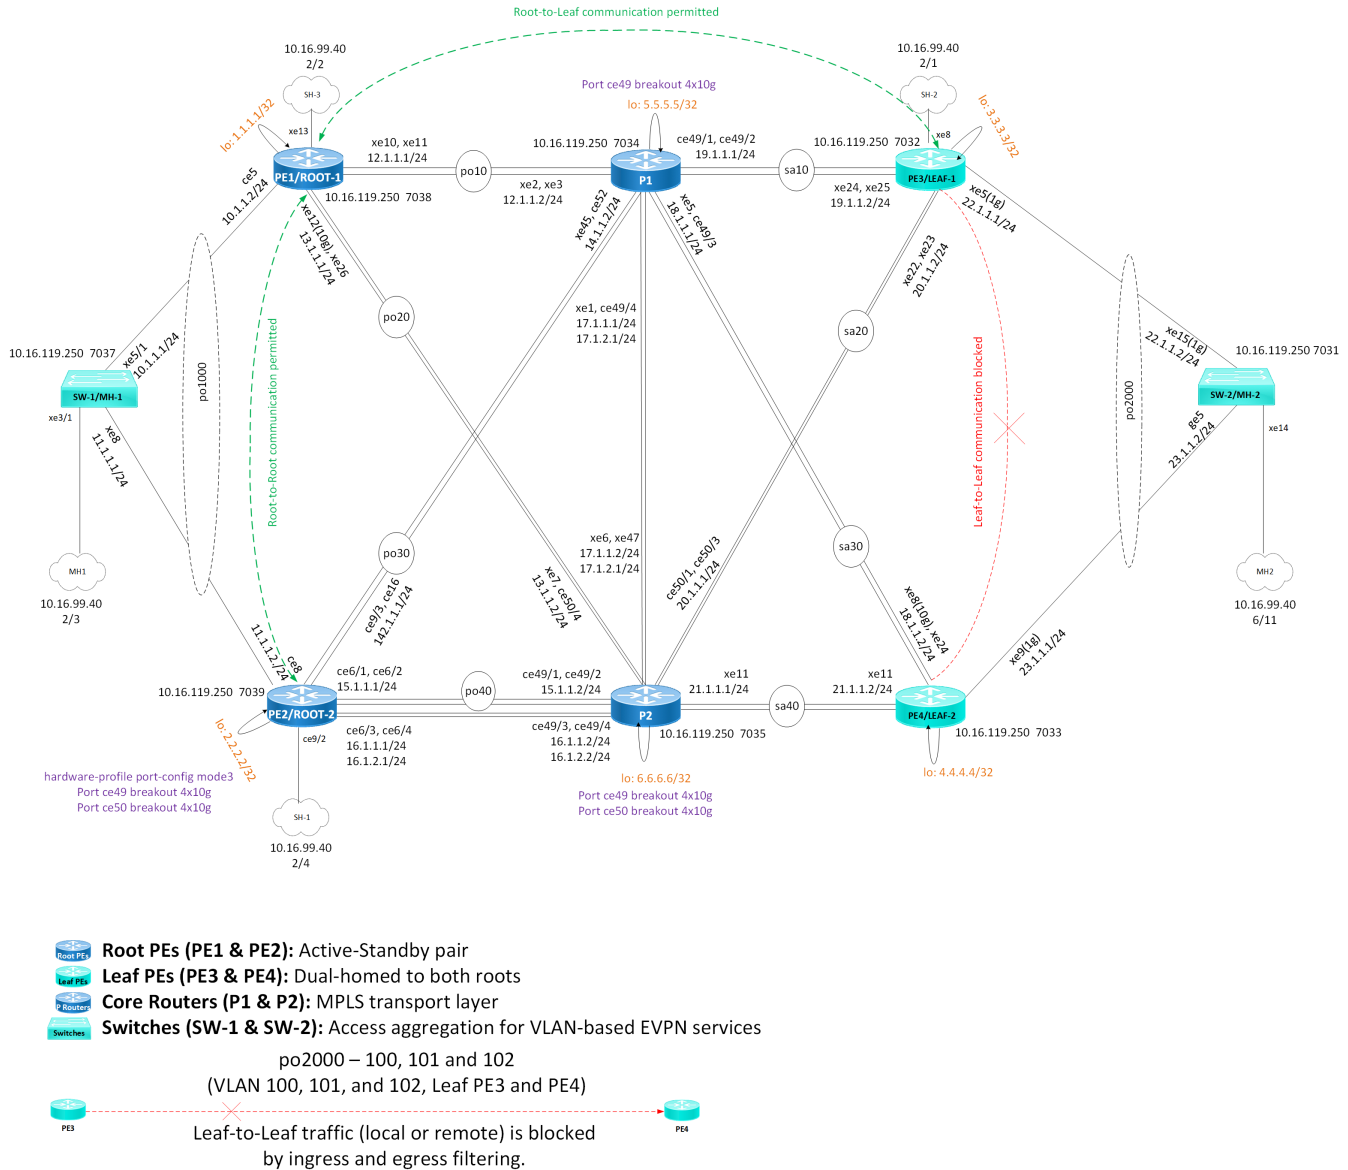
E-Tree Scenario 2 Configuration

This section describes a sample use case topology illustrating the E-Tree Scenario 2 deployment in an EVPN MPLS network. It explains enabling the scenario-2 on Provider Edge (PE) routers, its operation across the MPLS core, and validation steps to ensure proper functionality and traffic distribution.

Topology

| Node | Role or Function in EVPN E-Tree SC-2 |
|---------------|---|
| PE1 (Root) | Acts as a Root PE; originates and receives traffic from all Leaf and Root sites. Advertises RT-2 routes with Leaf indication for filtering. |
| PE2 (Root) | Redundant Root PE (active-standby pair with PE1). Provides Root-to-Root and Root-to-Leaf communication. Enforces ingress or egress filtering. |
| PE3 (Leaf) | Leaf PE (single-homed or multi-homed). Learns MACs from Leaf ACs and marks them with Leaf attributes. Cannot send directly to other Leaf ACs. |
| PE4 (Leaf) | Dual-homed Leaf PE (multi-homed with ESI). Sends BUM traffic with Leaf Label; remote PEs use label to drop traffic to Leaf sites. |
| P1 (Core) | MPLS core router transporting labeled EVPN traffic between PEs. Does not participate in E-Tree filtering but forwards based on labels. |
| P2 (Core) | MPLS core router providing redundancy and ECMP forwarding for EVPN or MPLS traffic. |
| SW-1 (Access) | Provides VLAN-based connectivity to Leaf subscribers; connects to Leaf ACs. Classified as Leaf-facing interface. |
| SW-2 (Access) | Provides VLAN-based connectivity to Root-facing services (e.g., data center). Classified as Root-facing interface. |

Figure 58. E-Tree Scenario 2



Note: Before configuration meet all [E-Tree Scenario 2 Prerequisites \(page 1890\)](#).

Enable E-Tree Scenario-2 Functionality

E-Tree Scenario-2 is enabled globally using the `evpn etree enable scenario-2` command. Activates E-Tree mode allowing both Root and Leaf sites on the same PE.

```
!
evpn etree enable scenario-2
!
hardware-profile filter evpn-mpls-mh enable
!
```

VPN ID and Leaf Role Mapping

- VPN IDs are defined for each E-Tree service, and access sub-interfaces are mapped with `map vpn-id <id> etree-leaf` where leaf identification is required. Associates sub-interfaces with E-Tree roles (Root or Leaf).
- Root sites communicate with all others; Leaf sites are isolated from each other. Configured using the `etree-leaf` parameter on access interfaces.

```
!
interface po1000
  switchport
  load-interval 30
  evpn multi-homed system-mac 0000.1212.1313
!
interface po1000.100 switchport
  encapsulation dot1q 100
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 100
!
interface po1000.101 switchport
  encapsulation dot1q 101
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 101 etree-leaf
!
interface po1000.102 switchport
  encapsulation dot1q 102
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 102 etree-leaf
!
```

E-Tree Scenario 2 Validation

After configuring the E-Tree Scenario 2 feature on all Provider Edge (PE) nodes, perform the following validation steps to confirm correct operation of ingress or egress filtering, Leaf Label advertisement, and BGP route exchange.



Notes:

- Repeat the same verification on all PEs (PE1–PE4) to confirm consistent Leaf Label allocation and DF or NON-DF synchronization.
- Ensure the `evpn etree enable scenario-2` and `map vpn-id <ID> etree-leaf` configurations are active before validation.

Verify DF and Non-DF Status on Multi-Homed PEs

- Displays per-EVI Designated Forwarder (DF) or Non-Designated Forwarder (Non-DF) status for multi-homed access circuits (ACs).
- In Scenario 2, DF or Non-DF roles ensure that only the DF PE forwards BUM traffic toward the shared Ethernet Segment Identifier (ESI), preventing duplication.
- Each VLAN or sub-interface under the same ESI should display alternating **DF** or **NON-DF** roles between MH-PEs.
- The below output confirms deterministic BUM handling across redundant PEs.

```

PE1#show evpn mpls
EVPN-MPLS Information
=====
Codes: NW - Network Port
       AC - Access Port
       (u) - Untagged

VPN-ID   EVI-Name   EVI-Type Type Interface ESI              VLAN   DF-Status
Src-Addr Dst-Addr
-----
100      ----      L2      NW    ----      ----              ----   ----
      1.1.1.1      3.3.3.3      ----
100      ----      L2      NW    ----      ----              ----   ----
      1.1.1.1      2.2.2.2      ----
100      ----      L2      NW    ----      ----              ----   ----
      1.1.1.1      4.4.4.4      ----
100      ----      --      AC    po1000.100 00:00:00:12:12:13:13:00:00:00 ----
      DF      ----      ----
101      ----      L2      NW    ----      ----              ----   ----
      1.1.1.1      3.3.3.3      ----
101      ----      L2      NW    ----      ----              ----   ----
      1.1.1.1      2.2.2.2      ----
101      ----      L2      NW    ----      ----              ----   ----
      1.1.1.1      4.4.4.4      ----
101      ----      --      AC    po1000.101 00:00:00:12:12:13:13:00:00:00 ----   NON-
      DF      ----      ----
102      ----      L2      NW    ----      ----              ----   ----
      1.1.1.1      3.3.3.3      ----
102      ----      L2      NW    ----      ----              ----   ----
      1.1.1.1      2.2.2.2      ----
102      ----      L2      NW    ----      ----              ----   ----
      1.1.1.1      4.4.4.4      ----
102      ----      --      AC    po1000.102 00:00:00:12:12:13:13:00:00:00 ----
      DF      ----      ----
Total number of entries are 12

```

Verify MPLS Tunnels Between PEs

- Ensures EVPN MPLS tunnels (per-EVI LSPs) between PEs are installed and operational.
- These tunnels provide underlay connectivity for E-Tree unicast and BUM forwarding.
- Tunnel status should be **Installed** for all peer PEs.
- Confirms label-switched connectivity between PEs, ensuring both Root and Leaf traffic can traverse the MPLS core.

```

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination   Status      Up/Down      Update      evpn-id      Local-
Leaf Remote-Leaf Ext-Color FAT
=====
1.1.1.1      3.3.3.3      Installed   00:17:08     00:17:08     102          ---
      ---      ---
1.1.1.1      3.3.3.3      Installed   00:17:08     00:17:08     101          ---
      ---      ---
1.1.1.1      3.3.3.3      Installed   00:17:07     00:17:07     100          ---
      ---      ---
1.1.1.1      2.2.2.2      Installed   00:17:24     00:17:24     102          ---
      ---      ---
1.1.1.1      2.2.2.2      Installed   00:17:24     00:17:24     101          ---
      ---      ---
1.1.1.1      2.2.2.2      Installed   00:17:24     00:17:24     100          ---
      ---      ---
1.1.1.1      4.4.4.4      Installed   00:17:21     00:17:21     102          ---
      ---      ---
1.1.1.1      4.4.4.4      Installed   00:17:21     00:17:21     101          ---
      ---      ---

```

```

1.1.1.1      4.4.4.4      Installed      00:17:21      00:17:21      100      ---
---          ---          ---

Total number of entries are 9

```

```
PE1#show evpn mpls tunnel summary
```

```
Total number of entries: 9 [Installed: 9, Resolved: 0, Unresolved: 0]
```

Verify E-Tree Leaf Label Association

- Displays mapping of access interfaces to their E-Tree roles (Root or Leaf) and corresponding Leaf Label assignment.
- Root ACs show **Root** with no Leaf Label.
- Leaf ACs show **Leaf** with a unique Leaf Label (per-PE).
- Confirms correct Leaf identification per access interface and validates local Leaf Label programming used for unicast or BUM filtering.

```
PE1#show evpn etree-leaf brief
```

| Leaf Interface | Ifindex | Vnid | Leaf status | label |
|-------------------|------------|------|----------------|-------|
| po1000.100 | 0x1f400064 | 100 | Root | ---- |
| po1000.101 | 0x1f400065 | 101 | Leaf | 16 |
| po1000.102 | 0x1f400066 | 102 | Leaf | 16 |

```
Total number of entries are 3
```

Verify Ethernet A-D per-ES (EAD) Route Advertisement

- Displays Ethernet Auto-Discovery (A-D) routes used for ESI discovery and Leaf Label distribution.
- In Scenario 2:
 - EAD routes with **ESI = 0** identify Leaf interfaces (single-homed or non-ESI).
 - EAD routes with **valid ESI values** represent multi-homed segments.
- Each Leaf AC advertises an A-D route with ESI = 0 and carries the assigned Leaf Label.
- Remote PEs receive these routes for BUM egress filtering.

```
RD [1.1.1.1:101]: ESI 0 VNID/LABEL 16 Nexthop 3.3.3.3
```

- Entries with ESI = 0 and Label = 16 correspond to Ethernet A-D per-ES (RT-1) routes carrying the Leaf Label.
- These routes confirm successful Leaf Label advertisement via BGP.
- The control-plane signaling for Leaf Label exchange is functioning, enabling:
 - Egress filtering for BUM traffic
 - Ingress filtering for unicast Leaf-to-Leaf communication

```
PE1#show bgp l2vpn evpn multihoming ethernet-ad-per-es
```

```
BGP table version is 7, local router ID is 1.1.1.1
```

```
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,
```

```
l - labeled, S Stale
```

Origin codes: i - IGP, e - EGP, ? - incomplete
 Description : Ext-Color - Extended community color

RD[1.1.1.1:100] VRF[vrf100]:

| IP | ESI | Peer | Encap | Eth-Tag Flags | VNID/LABEL | Nexthop | |
|----|-------------------------------|------|-------|------------------|------------|---------|--------------|
| * | | | | | | | |
| i | 00:00:00:12:12:13:13:00:00:00 | | | 4294967295 | 304 | 2.2.2.2 | 2.2.2.2 MPLS |
| * | | | | | | | |
| i | 00:00:00:14:14:15:15:00:00:00 | | | 4294967295 | 304 | 4.4.4.4 | 4.4.4.4 MPLS |
| * | | | | | | | |
| i | 00:00:00:14:14:15:15:00:00:00 | | | 4294967295 | 304 | 3.3.3.3 | 3.3.3.3 MPLS |

RD[1.1.1.1:101] VRF[vrf101]:

| IP | ESI | Peer | Encap | Eth-Tag Flags | VNID/LABEL | Nexthop | |
|----|-------------------------------|------|-------|------------------|------------|---------|--------------|
| * | | | | | | | |
| i | 0 | | | 4294967295 | 16 | 3.3.3.3 | 3.3.3.3 MPLS |
| * | | | | | | | |
| i | 0 | | | 4294967295 | 16 | 4.4.4.4 | 4.4.4.4 MPLS |
| * | | | | | | | |
| i | 0 | | | 4294967295 | 16 | 2.2.2.2 | 2.2.2.2 MPLS |
| * | | | | | | | |
| i | 00:00:00:12:12:13:13:00:00:00 | | | 4294967295 | 304 | 2.2.2.2 | 2.2.2.2 MPLS |
| * | | | | | | | |
| i | 00:00:00:14:14:15:15:00:00:00 | | | 4294967295 | 304 | 4.4.4.4 | 4.4.4.4 MPLS |
| * | | | | | | | |
| i | 00:00:00:14:14:15:15:00:00:00 | | | 4294967295 | 304 | 3.3.3.3 | 3.3.3.3 MPLS |

RD[1.1.1.1:102] VRF[vrf102]:

| IP | ESI | Peer | Encap | Eth-Tag Flags | VNID/LABEL | Nexthop | |
|----|-------------------------------|------|-------|------------------|------------|---------|--------------|
| * | | | | | | | |
| i | 0 | | | 4294967295 | 16 | 3.3.3.3 | 3.3.3.3 MPLS |
| * | | | | | | | |
| i | 0 | | | 4294967295 | 16 | 4.4.4.4 | 4.4.4.4 MPLS |
| * | | | | | | | |
| i | 0 | | | 4294967295 | 16 | 2.2.2.2 | 2.2.2.2 MPLS |
| * | | | | | | | |
| i | 00:00:00:12:12:13:13:00:00:00 | | | 4294967295 | 304 | 2.2.2.2 | 2.2.2.2 MPLS |
| * | | | | | | | |
| i | 00:00:00:14:14:15:15:00:00:00 | | | 4294967295 | 304 | 4.4.4.4 | 4.4.4.4 MPLS |
| * | | | | | | | |
| i | 00:00:00:14:14:15:15:00:00:00 | | | 4294967295 | 304 | 3.3.3.3 | 3.3.3.3 MPLS |

RD[1.1.1.1:64512] VRF[evpn-gvrf-1]:

| IP | ESI | Peer | Encap | Eth-Tag Flags | VNID/LABEL | Nexthop | |
|----|-------------------------------|------|-------|------------------|------------|---------|-------|
| *> | 0 | | | 4294967295 | 16 | 1.1.1.1 | ----- |
| | | MPLS | | | | | |
| *> | 00:00:00:12:12:13:13:00:00:00 | | | 4294967295 | 304 | 1.1.1.1 | ----- |
| | | MPLS | | | | | |

RD[2.2.2.2:64512]

| IP | ESI | Peer | Encap | Eth-Tag Flags | VNID/LABEL | Nexthop | |
|-----|-------------------------------|------|-------|------------------|------------|---------|------------|
| *>i | 0 | | | 4294967295 | 16 | 2.2.2.2 | 2.2.2.2 MP |
| LS | | | | | | | |
| *>i | 00:00:00:12:12:13:13:00:00:00 | | | 4294967295 | 304 | 2.2.2.2 | 2.2.2.2 MP |
| LS | | | | | | | |

RD[3.3.3.3:64512]

| IP | ESI | Peer | Encap | Eth-Tag Flags | VNID/LABEL | Nexthop | |
|-----|-------------------------------|------|-------|------------------|------------|---------|------------|
| *>i | 0 | | | 4294967295 | 16 | 3.3.3.3 | 3.3.3.3 MP |
| LS | | | | | | | |
| *>i | 00:00:00:14:14:15:15:00:00:00 | | | 4294967295 | 304 | 3.3.3.3 | 3.3.3.3 MP |
| LS | | | | | | | |

```
RD[4.4.4.4:64512]
      ESI              Eth-Tag      VNID/LABEL      Nexthop
IP      Peer          Encap        Flags
*>i  0                4294967295  16              4.4.4.4          4.4.4.4          MP
LS
*>i  00:00:00:14:14:15:15:00:00:00  4294967295  304             4.4.4.4          4.4.4.4          MP
LS
```

E-Tree Scenario 2 Network Topology Running Configurations

Here are the snippet configurations for all nodes in the given network topology [Figure 58](#).

Root PEs

PE1

```
PE1#show running-config

!
hardware-profile filter ingress-ipv6-ext-subif enable
hardware-profile filter ingress-ipv4-subif enable
hardware-profile filter egress-ipv4-ext enable
hardware-profile filter egress-ipv6 enable
hardware-profile filter evpn-mpls-mh enable
hardware-profile filter twamp-ipv4 enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
!
hostname PE1-7038
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature ssh vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
feature rsyslog
logging remote server 10.16.100.20 5 port 1514 vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
evpn mpls enable
!
evpn esi hold-time 200
!
evpn etree enable scenario-2
!
evpn mpls multihoming enable
!
ip vrf management
!
mac vrf vrf100
rd 1.1.1.1:100
```

```
    route-target both 65010:100
!
mac vrf vrf101
  rd 1.1.1.1:101
  route-target both 65010:101
!
mac vrf vrf102
  rd 1.1.1.1:102
  route-target both 65010:102
!
evpn mpls vtep-ip-global 1.1.1.1
!
evpn mpls id 100
  host-reachability-protocol evpn-bgp vrf100
!
evpn mpls id 101
  host-reachability-protocol evpn-bgp vrf101
!
evpn mpls id 102
  host-reachability-protocol evpn-bgp vrf102
!
router ldp
  graceful-restart full
  targeted-peer ipv4 2.2.2.2
    exit-targeted-peer-mode
  targeted-peer ipv4 3.3.3.3
    exit-targeted-peer-mode
  targeted-peer ipv4 4.4.4.4
    exit-targeted-peer-mode
  transport-address ipv4 1.1.1.1
!
router rsvp
!
interface po10
  description connected-to-p1
  load-interval 30
  ip address 12.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface po20
  description connected-to-p2
  load-interval 30
  ip address 13.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface po1000
  switchport
  load-interval 30
  evpn multi-homed system-mac 0000.1212.1313
!
interface po1000.100 switchport
  encapsulation dot1q 100
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 100
!
interface po1000.101 switchport
  encapsulation dot1q 101
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
```



```
    map vpn-id 101 etree-leaf
!
interface po1000.102 switchport
 encapsulation dot1q 102
 rewrite push 0x8100 4090
 load-interval 30
 access-if-evpn
  map vpn-id 102 etree-leaf
!
interface ce5
 channel-group 1000 mode active
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
 ip router isis ISIS-IGP
 enable-rsvp
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe10
 channel-group 10 mode active
!
interface xe11
 channel-group 10 mode active
!
interface xe12
 speed 10g
 channel-group 20 mode active
!
interface xe26
 channel-group 20 mode active
!
exit
!
router ospf 100
 ospf router-id 1.1.1.1
 bfd all-interfaces
 network 1.1.1.1/32 area 0.0.0.0
 network 12.1.1.0/24 area 0.0.0.0
 network 13.1.1.0/24 area 0.0.0.0
!
router bgp 65010
 neighbor EVPN peer-group
 neighbor EVPN remote-as 65010
 neighbor EVPN update-source lo
 neighbor EVPN advertisement-interval 0
 neighbor EVPN fall-over bfd multihop
 neighbor 2.2.2.2 peer-group EVPN
 neighbor 3.3.3.3 peer-group EVPN
 neighbor 4.4.4.4 peer-group EVPN
!
 address-family l2vpn evpn
  neighbor EVPN activate
 exit-address-family
!
exit
!
rsvp-path PE1-PE3 mpls
 13.1.1.2 strict
```

```

17.1.1.1 strict
19.1.1.2 strict
!
rsvp-path PE1-PE4 mpls
12.1.1.2 strict
17.1.1.2 strict
21.1.1.2 strict
!
rsvp-trunk PE1-PE3 ipv4
primary fast-reroute protection one-to-one
primary path PE1-PE3
to 3.3.3.3
!
rsvp-trunk PE1-PE4 ipv4
primary fast-reroute protection facility
primary path PE1-PE4
to 4.4.4.4
!
rsvp-trunk PE1-PE2 ipv4
to 2.2.2.2
!
line console 0
exec-timeout 0
!
!
end
!

```

PE2

```

PE2#show running-config

!
hardware-profile filter ingress-ipv6-ext-subif enable
hardware-profile filter ingress-ipv4-subif enable
hardware-profile filter egress-ipv4-ext enable
hardware-profile filter egress-ipv6 enable
hardware-profile filter evpn-mpls-mh enable
hardware-profile filter twamp-ipv4 enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
hardware-profile port-config mode3
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE2-7039
port ce9 breakout 4X10g
port ce6 breakout 4X10g
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature ssh vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
lldp run
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt management-address
!
evpn mpls enable
!
evpn esi hold-time 200

```

```
!  
evpn etree enable scenario-2  
!  
evpn mpls multihoming enable  
!  
ip vrf management  
!  
mac vrf vrf100  
  rd 2.2.2.2:100  
  route-target both 65010:100  
!  
mac vrf vrf101  
  rd 2.2.2.2:101  
  route-target both 65010:101  
!  
mac vrf vrf102  
  rd 2.2.2.2:102  
  route-target both 65010:102  
!  
evpn mpls vtep-ip-global 2.2.2.2  
!  
evpn mpls id 100  
  host-reachability-protocol evpn-bgp vrf100  
!  
evpn mpls id 101  
  host-reachability-protocol evpn-bgp vrf101  
!  
evpn mpls id 102  
  host-reachability-protocol evpn-bgp vrf102  
!  
router ldp  
  graceful-restart full  
  targeted-peer ipv4 1.1.1.1  
    exit-targeted-peer-mode  
  targeted-peer ipv4 3.3.3.3  
    exit-targeted-peer-mode  
  targeted-peer ipv4 4.4.4.4  
    exit-targeted-peer-mode  
  transport-address ipv4 2.2.2.2  
!  
router rsvp  
!  
interface po30  
  description connected-to-pl  
  load-interval 30  
  ip address 14.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface po40  
  load-interval 30  
  ip address 15.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface po1000  
  switchport  
  load-interval 30  
  evpn multi-homed system-mac 0000.1212.1313  
!  
interface po1000.100 switchport  
  encapsulation dot1q 100  
  rewrite push 0x8100 4090  
  load-interval 30
```

```
access-if-evpn
  map vpn-id 100
!
interface po1000.101 switchport
  encapsulation dot1q 101
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
    map vpn-id 101 etree-leaf
!
interface po1000.102 switchport
  encapsulation dot1q 102
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
    map vpn-id 102 etree-leaf
!
interface ce6/1
  channel-group 40 mode active
!
interface ce6/2
  channel-group 40 mode active
!
interface ce6/3
  description connected-to-p2
  load-interval 30
  ip address 16.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce6/4
  description connected-to-p2
  load-interval 30
  ip address 16.1.2.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce7
!
interface ce8
  channel-group 1000 mode active
!
interface ce9/1
!
interface ce9/2
!
interface ce9/2.101 switchport
  encapsulation dot1q 101
  load-interval 30
  access-if-evpn
    map vpn-id 101 etree-leaf
!
interface ce9/2.102 switchport
  encapsulation dot1q 102
  load-interval 30
  access-if-evpn
    map vpn-id 102 etree-leaf
!
interface ce9/3
  channel-group 30 mode active
!
interface ce16
  channel-group 30 mode active
!
```

```
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 2.2.2.2/32 secondary
 ipv6 address ::1/128
 ip router isis ISIS-IGP
 enable-rsvp
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe0
!
interface xe1
 speed 10g
!
exit
!
router ospf 100
 ospf router-id 2.2.2.2
 bfd all-interfaces
 network 2.2.2.2/32 area 0.0.0.0
 network 14.1.1.0/24 area 0.0.0.0
 network 15.1.1.0/24 area 0.0.0.0
 network 16.1.1.0/24 area 0.0.0.0
 network 16.1.2.0/24 area 0.0.0.0
!
router bgp 65010
 neighbor EVPN peer-group
 neighbor EVPN remote-as 65010
 neighbor EVPN update-source lo
 neighbor EVPN advertisement-interval 0
 neighbor EVPN fall-over bfd multihop
 neighbor 1.1.1.1 peer-group EVPN
 neighbor 3.3.3.3 peer-group EVPN
 neighbor 4.4.4.4 peer-group EVPN
!
 address-family l2vpn evpn
 neighbor EVPN activate
 exit-address-family
!
exit
!
rsvp-path PE2-PE3 mpls
 15.1.1.2 loose
 17.1.1.1 loose
 19.1.1.2 loose
!
rsvp-path PE2-PE4 mpls
 14.1.1.2 loose
 17.1.1.2 loose
 21.1.1.2 loose
!
rsvp-trunk PE2-PE3 ipv4
 primary fast-reroute protection one-to-one
 primary path PE2-PE3
 to 3.3.3.3
!
rsvp-trunk PE2-PE4 ipv4
 primary fast-reroute protection facility
 primary path PE2-PE4
 to 4.4.4.4
!
```

```
rsvp-trunk PE2-PE1 ipv4
to 1.1.1.1
!
line console 0
exec-timeout 0
!
!
end

!
```

Leaf PEs

PE3

```
PE3#show running-config
!
hardware-profile filter ingress-ipv6-ext-subif enable
hardware-profile filter ingress-ipv4-subif enable
hardware-profile filter egress-ipv4-ext enable
hardware-profile filter egress-ipv6 enable
hardware-profile filter evpn-mpls-mh enable
hardware-profile filter twamp-ipv4 enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE3-7032
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature ssh vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
lldp run
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt management-address
!
evpn mpls enable
!
evpn etree enable scenario-2
!
evpn mpls multihoming enable
!
ip vrf management
!
mac vrf vrf100
rd 3.3.3.3:100
route-target both 65010:100
!
mac vrf vrf101
rd 3.3.3.3:101
route-target both 65010:101
!
mac vrf vrf102
rd 3.3.3.3:102
route-target both 65010:102
evpn mpls vtep-ip-global 3.3.3.3
!
```

```
evpn mpls id 100
  host-reachability-protocol evpn-bgp vrf100
!
evpn mpls id 101
  host-reachability-protocol evpn-bgp vrf101
!
evpn mpls id 102
  host-reachability-protocol evpn-bgp vrf102
!
router ldp
  graceful-restart full
  targeted-peer ipv4 1.1.1.1
    exit-targeted-peer-mode
  targeted-peer ipv4 2.2.2.2
    exit-targeted-peer-mode
  targeted-peer ipv4 4.4.4.4
    exit-targeted-peer-mode
  transport-address ipv4 3.3.3.3
!
router rsvp
!
interface po2000
  switchport
  load-interval 30
  evpn multi-homed system-mac 0000.1414.1515
!
interface po2000.100 switchport
  encapsulation dot1q 100
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 100
!
interface po2000.101 switchport
  encapsulation dot1q 101
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 101 etree-leaf
!
interface po2000.102 switchport
  encapsulation dot1q 102
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 102 etree-leaf
!
interface sa10
  description connected-to-p1
  load-interval 30
  ip address 19.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface sa20
  description connected-to-p2
  load-interval 30
  ip address 20.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
```

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 3.3.3.3/32 secondary  
  ipv6 address ::1/128  
  ip router isis ISIS-IGP  
  enable-rsvp  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xe4  
!  
interface xe5  
  speed lg  
  channel-group 2000 mode active  
!  
interface xe8.101 switchport  
  encapsulation dot1q 101  
  load-interval 30  
  access-if-evpn  
  map vpn-id 101  
!  
interface xe22  
  static-channel-group 20  
!  
interface xe23  
  static-channel-group 20  
!  
interface xe24  
  static-channel-group 10  
!  
interface xe25  
  static-channel-group 10  
!  
exit  
!  
router ospf 100  
  ospf router-id 3.3.3.3  
  bfd all-interfaces  
  network 3.3.3.3/32 area 0.0.0.0  
  network 19.1.1.0/24 area 0.0.0.0  
  network 20.1.1.0/24 area 0.0.0.0  
!  
router bgp 65010  
  neighbor EVPN peer-group  
  neighbor EVPN remote-as 65010  
  neighbor EVPN update-source lo  
  neighbor EVPN advertisement-interval 0  
  neighbor EVPN fall-over bfd multihop  
  neighbor 1.1.1.1 peer-group EVPN  
  neighbor 2.2.2.2 peer-group EVPN  
  neighbor 4.4.4.4 peer-group EVPN  
!  
  address-family l2vpn evpn  
  neighbor EVPN activate  
  exit-address-family  
!  
exit  
!  
  rsvp-path PE3-PE1 mpls  
    20.1.1.1 strict  
    17.1.1.1 strict  
    12.1.1.1 strict  
  !  
  rsvp-path PE3-PE2 mpls
```



```

19.1.1.1 strict
17.1.1.2 strict
15.1.1.1 strict
!
rsvp-trunk PE3-PE1 ipv4
primary fast-reroute protection facility
primary path PE3-PE1
to 1.1.1.1
!
rsvp-trunk PE3-PE2 ipv4
primary fast-reroute protection one-to-one
primary path PE3-PE2
to 2.2.2.2
!
rsvp-trunk PE3-PE4 ipv4
to 4.4.4.4
!
line console 0
exec-timeout 0
!
!
end
!

```

PE4

```

PE4#show running-config

!
hardware-profile filter ingress-ipv6-ext-subif enable
hardware-profile filter ingress-ipv4-subif enable
hardware-profile filter egress-ipv4-ext enable
hardware-profile filter egress-ipv6 enable
hardware-profile filter evpn-mpls-mh enable
hardware-profile filter twamp-ipv4 enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE4-7033
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
evpn mpls arp-nd refresh-timer 120
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt management-address
!
evpn mpls enable
!
evpn esi hold-time 200
!
evpn etree enable scenario-2
!
evpn mpls multihoming enable
!
ip vrf management
!
mac vrf vrf100

```

```
rd 4.4.4.4:100
route-target both 65010:100
!
mac vrf vrf101
rd 4.4.4.4:101
route-target both 65010:101
!
mac vrf vrf102
rd 4.4.4.4:102
route-target both 65010:102
!
evpn mpls vtep-ip-global 4.4.4.4
!
evpn mpls mac-ageing-time 360
!
evpn mpls id 100
host-reachability-protocol evpn-bgp vrf100
!
evpn mpls id 101
host-reachability-protocol evpn-bgp vrf101
!
evpn mpls id 102
host-reachability-protocol evpn-bgp vrf102
!
router ldp
graceful-restart full
targeted-peer ipv4 1.1.1.1
exit-targeted-peer-mode
targeted-peer ipv4 2.2.2.2
exit-targeted-peer-mode
targeted-peer ipv4 3.3.3.3
exit-targeted-peer-mode
transport-address ipv4 4.4.4.4
!
router rsvp
!
interface po2000
switchport
load-interval 30
evpn multi-homed system-mac 0000.1414.1515
!
interface po2000.100 switchport
encapsulation dot1q 100
rewrite push 0x8100 4090
load-interval 30
access-if-evpn
map vpn-id 100
!
interface po2000.101 switchport
encapsulation dot1q 101
rewrite push 0x8100 4090
load-interval 30
access-if-evpn
map vpn-id 101 etree-leaf
!
interface po2000.102 switchport
encapsulation dot1q 102
rewrite push 0x8100 4090
load-interval 30
access-if-evpn
map vpn-id 102 etree-leaf
!
interface sa30
description connected-to-p1
load-interval 30
ip address 18.1.1.2/24
mtu 9216
label-switching
```

```
enable-ldp ipv4
enable-rsvp
!
interface sa40
description connected-to-p2
load-interval 30
ip address 21.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface cd0
!
interface cd1
!
interface ce2
!
interface ce3
!
interface eth0
ip vrf forwarding management
ip address dhcp
!
interface lo
ip address 127.0.0.1/8
ip address 4.4.4.4/32 secondary
ipv6 address ::1/128
ip router isis ISIS-IGP
enable-rsvp
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface xe6
speed 10g
!
interface xe7
!
interface xe8
speed 10g
static-channel-group 30
!
interface xe9
speed 1g
channel-group 2000 mode active
!
interface xe10
!
interface xe11
static-channel-group 40
!
interface xe24
static-channel-group 30
!
exit
!
router ospf 100
ospf router-id 4.4.4.4
bfd all-interfaces
network 4.4.4.4/32 area 0.0.0.0
network 18.1.1.0/24 area 0.0.0.0
network 21.1.1.0/24 area 0.0.0.0
!
router bgp 65010
neighbor EVPN peer-group
```

```

neighbor EVPN remote-as 65010
neighbor EVPN update-source lo
neighbor EVPN advertisement-interval 0
neighbor EVPN fall-over bfd multihop
neighbor 1.1.1.1 peer-group EVPN
neighbor 2.2.2.2 peer-group EVPN
neighbor 3.3.3.3 peer-group EVPN
!
address-family l2vpn evpn
neighbor EVPN activate
exit-address-family
!
exit
!
!
rsvp-path PE4-PE1 mpls
21.1.1.1 strict
13.1.1.1 strict
!
rsvp-path PE4-PE2 mpls
18.1.1.1 strict
14.1.1.1 strict
!
rsvp-trunk PE4-PE1 ipv4
primary fast-reroute protection one-to-one
primary path PE4-PE1
to 1.1.1.1
!
rsvp-trunk PE4-PE2 ipv4
primary fast-reroute protection facility
primary path PE4-PE2
to 2.2.2.2
!
rsvp-trunk PE4-PE3 ipv4
to 3.3.3.3
!
line console 0
exec-timeout 0
!
!
end
!

```

Core Routers

P1

```

P1#show running-config
!
hardware-profile statistics ingress-acl enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname P1-7034
port ce49 breakout 4X10g
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt system-name

```

```
lldp tlv-select basic-mgmt management-address
!
ip vrf management
!
router ldp
  transport-address ipv4 5.5.5.5
!
router rsvp
!
interface po10
  load-interval 30
  ip address 12.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface po30
  description connected-to-pe2
  load-interval 30
  ip address 14.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface sa10
  description connected-to-pe3
  load-interval 30
  ip address 19.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface sa30
  description connected-to-pe4
  load-interval 30
  ip address 18.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce49/1
  static-channel-group 10
!
interface ce49/2
  static-channel-group 10
!
interface ce49/3
  static-channel-group 30
!
interface ce49/4
  description connected-to-p2
  load-interval 30
  ip address 17.1.2.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce52
  channel-group 30 mode active
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
```

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 5.5.5.5/32 secondary  
  ipv6 address ::1/128  
  ip router isis ISIS-IGP  
  enable-rsvp  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xe1  
  description connected-to-p2  
  load-interval 30  
  ip address 17.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface xe2  
  channel-group 10 mode active  
!  
interface xe3  
  channel-group 10 mode active  
!  
interface xe4  
!  
interface xe5  
  static-channel-group 30  
!  
interface xe45  
  channel-group 30 mode active  
!  
  exit  
!  
router ospf 100  
  ospf router-id 5.5.5.5  
  bfd all-interfaces  
  network 5.5.5.5/32 area 0.0.0.0  
  network 12.1.1.0/24 area 0.0.0.0  
  network 14.1.1.0/24 area 0.0.0.0  
  network 17.1.1.0/24 area 0.0.0.0  
  network 17.1.2.0/24 area 0.0.0.0  
  network 18.1.1.0/24 area 0.0.0.0  
  network 19.1.1.0/24 area 0.0.0.0  
!  
line console 0  
  exec-timeout 0  
!  
!  
end  
!
```

P2

```
P2#show running-config  
!  
hardware-profile statistics ingress-acl enable  
!  
bfd interval 3 minrx 3 multiplier 3  
!  
qos enable
```

```
!  
hostname P2-7035  
port ce49 breakout 4X10g  
port ce50 breakout 4X10g  
ip name-server vrf management 10.16.10.23  
tfo Disable  
errdisable cause stp-bpdu-guard  
feature dns relay  
ip dns relay  
ipv6 dns relay  
feature rsyslog  
logging remote server 10.16.100.20 5 port 1514 vrf management  
lldp run  
lldp tlv-select basic-mgmt port-description  
lldp tlv-select basic-mgmt system-name  
lldp tlv-select basic-mgmt system-capabilities  
lldp tlv-select basic-mgmt system-description  
lldp tlv-select basic-mgmt management-address  
lldp notification-interval 1000  
fault-management enable  
!  
ip vrf management  
!  
router ldp  
  transport-address ipv4 6.6.6.6  
!  
router rsvp  
!  
interface po20  
  description connected-to-pe1  
  load-interval 30  
  ip address 13.1.1.2/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface po40  
  description connected-to-pe2  
  load-interval 30  
  ip address 15.1.1.2/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface sa20  
  description connected-to-pe3  
  load-interval 30  
  ip address 20.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface sa40  
  description connected-to-pe4  
  load-interval 30  
  ip address 21.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface ce49/1  
  channel-group 40 mode active  
!  
interface ce49/2
```

```
channel-group 40 mode active
!
interface ce49/3
description connected-to-pe2
load-interval 30
ip address 16.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface ce49/4
description connected-to-pe2
load-interval 30
ip address 16.1.2.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface ce50/1
static-channel-group 20
!
interface ce50/2
!
interface ce50/3
static-channel-group 20
!
interface ce50/4
channel-group 20 mode active
!
interface eth0
ip vrf forwarding management
ip address dhcp
!
interface lo
ip address 127.0.0.1/8
ip address 6.6.6.6/32 secondary
ipv6 address ::1/128
ip router isis ISIS-IGP
enable-rsvp
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface xe6
description connected-to-pl
load-interval 30
ip address 17.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface xe7
channel-group 20 mode active
!

interface xell
static-channel-group 40
!
interface xe20
label-switching
!
interface xe47
description connected-to-pl
```



```
load-interval 30
ip address 17.1.2.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface xe48
!
exit
!
router ospf 100
ospf router-id 6.6.6.6
bfd all-interfaces
network 6.6.6.6/32 area 0.0.0.0
network 13.1.1.0/24 area 0.0.0.0
network 15.1.1.0/24 area 0.0.0.0
network 16.1.1.0/24 area 0.0.0.0
network 16.1.2.0/24 area 0.0.0.0
network 17.1.1.0/24 area 0.0.0.0
network 17.1.2.0/24 area 0.0.0.0
network 20.1.1.0/24 area 0.0.0.0
network 21.1.1.0/24 area 0.0.0.0
!
line console 0
exec-timeout 0
!
!
end
!
```

Access Switches

SW-1

```
SW1#show running-config
!
ip vrf management
!
hostname SW1-7037
ip domain-lookup
bridge 1 protocol ieee vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
data-center-bridging enable bridge 1
feature telnet vrf
no feature telnet
feature ssh vrf
no feature ssh
snmp-server enable snmp vrf
snmp-server view all .1 included vrf
feature ntp vrf
ntp enable vrf
feature rsyslog vrf
lldp run
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-description
!
vlan database
vlan-reservation 4000-4062
vlan 2-1000 bridge 1 state enable
!
interface po1000
switchport
```

```

bridge-group 1 spanning-tree disable
switchport mode trunk
switchport trunk allowed vlan all
!
interface lo
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface xe1/1
!
interface xe1/2
!
interface xe1/3
!
interface xe1/4
!
interface xe2
!
interface xe3/1
port breakout enable 4X10g
switchport
bridge-group 1 spanning-tree disable
switchport mode trunk
switchport trunk allowed vlan all
!
interface xe3/2
!
interface xe5/1
channel-group 1000 mode active
!
interface xe8
channel-group 1000 mode active
!
no mac-address-table learning bridge 1 interface po1000
no mac-address-table learning bridge 1 interface xe3/1
!
line console 0
exec-timeout 0
!
!
end

```

E-Tree Scenario 2 Implementation Examples

In a real-world enterprise WAN or broadband service scenario, multiple branch offices connect to regional PE devices. Using OcNOS EVPN E-Tree Scenario 2, the provider ensures that branches (Leaf sites) cannot communicate directly with each other, enforcing strict isolation. Data centers (Spine sites) continue to communicate with all branches and other Root sites. EVPN E-Tree Scenario 2 uses Leaf Labels, ingress unicast filtering, and egress BUM filtering to enforce these rules efficiently, even on PE devices hosting both Root and Leaf sites. This design supports both single-homed and multi-homed branches, providing a secure, scalable, and high-performance solution for multi-tenant or enterprise deployments.

E-Tree Scenario 2 Troubleshooting

- Missing RT-2 route attributes: Verify per-AC or EVI configuration.

- Leaf Label not programmed: Check route reception and hardware resources.
- Unexpected Leaf-to-Leaf forwarding: Confirm Leaf roles and ingress filtering rules.
- BUM flooding to Leaf ACs: Check Leaf Label advertisement in Ethernet A-D routes.
- CLI not active: Ensure `evpn etree enable scenario-2` command is configured.

E-Tree Scenario 2 Revised Commands

- Introduced the scenario-2 parameter in the [evpn etree \(page 1920\)](#) command to enable EVPN MPLS E-Tree SC-2, which configures the device to support both Root and Leaf sites on the same PE device.
- Added `etree-leaf` parameter in the [map vpn-id \(page 1921\)](#) to enable configuring the access interface (AC) or sub-interface as a Leaf in an EVPN E-Tree deployment.

evpn etree

Use this command to enable E-Tree functionality within the EVPN configuration. The `evpn etree enable` option ([EVPN MPLS E-Tree Scenario 1 \(page 1859\)](#)) configures the device to function as either a Leaf or a Root site per EVPN Instance (EVI), as defined in RFC-8317. The `scenario-2` option ([EVPN MPLS E-Tree Scenario 2 \(page 1888\)](#)) configures the device to support both Root and Leaf sites on the same Provider Edge (PE) device, enforcing traffic filtering rules to prevent Leaf-to-Leaf communication while allowing Root-to-Root and Root-to-Leaf traffic.



Note: If Scenario 1 is configured, it must be removed before enabling Scenario 2.

Command Syntax

```
evpn etree enable (scenario-2|)
no evpn etree enable (scenario-2|)
```

Parameters

scenario-2

Enables EVPN MPLS E-Tree Scenario 2, which configures the device to support both Root and Leaf sites on the same PE device.

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.1. Introduced the `scenario-2` parameter in OcNOS version 7.0.0.

Example

The following example illustrates how to activate E-Tree Scenario 1 functionality for EVPN, which configures the device as either a Leaf or a Root site per EVI.

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable

OcNOS(config)#no evpn etree enable
```

The following example illustrates how to activate E-Tree Scenario 2 functionality for EVPN.

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable scenario-2
```

map vpn-id

Use this command to map a sub-interface to a tenant.

Use the no form of this command to remove the tenant.

Command Syntax

```
map vpn-id <1-16777215> (etree-leaf|)  
no map vpn-id <1-16777215>
```

Parameters

vpn-id <1-16777215>

Specifies the EVPN Instance (EVI) ID (VNID) to which the interface belongs.

etree-leaf

(Optional) Enables configuring the access interface (AC) or sub-interface as a Leaf in an EVPN E-Tree deployment. This allows the PE to enforce E-Tree traffic rules, including Leaf-to-Leaf traffic blocking and Root-to-Leaf communication.

Command Mode

Access interface EVPN mode

Applicability

Introduced in OcNOS version 3.0. Added `etree-leaf` parameter in OcNOS version 7.0.0.

Examples

1. Configure a regular EVPN interface (no E-Tree Leaf):

```
#configure terminal  
(config)#interface xel.1 switchport  
(config-if)#access-if-evpn  
(config-access-if)#map vpn-id 1  
(config-access-if)#end
```

2. Configure an interface as an E-Tree Leaf AC:

```
#configure terminal  
(config)#evpn etree enable scenario-2  
(config)#evpn mpls id 203 etree-leaf  
(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf103  
(config-evpn-mpls)#exit  
  
(config)#interface xel.1 switchport  
(config-if)#access-if-evpn  
(config-access-if)#map vpn-id 3 etree-leaf  
(config-access-if)#end
```

E-Tree Scenario 2 Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms or Acronym | Description |
|---|--|
| Attachment Circuit (AC) | Physical or logical interface (e.g., sub-interface, port-channel) connecting a customer VLAN or segment to a PE router. |
| Broadcast, Unknown-unicast, Multicast (BUM) | Collective term for Layer 2 traffic types that require flooding or replication across EVPN sites. |
| Designated Forwarder (DF) | PE router elected to forward BUM traffic on a multi-homed Ethernet Segment (ESI). Prevents duplication. |
| Ethernet Auto-Discovery (EAD) | EVPN route type used for signaling Ethernet Segment (ES) membership and local Leaf Label advertisement. |
| E-Tree | Ethernet VPN service model that defines hierarchical connectivity between Root and Leaf sites, restricting Leaf-to-Leaf communication. |
| Ethernet Segment (ES) | Group of physical links connecting a customer device (CE or switch) redundantly to two or more PEs, identified by an ESI. |
| Ethernet Segment Identifier (ESI) | 10-byte unique identifier representing a multi-homed access link in EVPN. |
| Ethernet VPN (EVPN) | BGP-based VPN technology providing Layer 2 services over MPLS or IP networks. |
| EVPN Instance (EVI) | Logical service instance that defines the scope of MAC learning and forwarding for a given EVPN. |
| Interior Gateway Protocol (IGP) | Routing protocol (e.g., IS-IS or OSPF) used to exchange underlay routes among core and PE routers. |
| Label-Switching | MPLS data-plane forwarding method that uses labels to route packets efficiently through the network. |
| Leaf Label | Locally assigned label identifying a Leaf access circuit; used to enforce Leaf-to-Root communication only. |
| Label Distribution Protocol (LDP) | MPLS control protocol that distributes label bindings for IGP prefixes. |
| Multiprotocol Label Switching (MPLS) | Transport technology that uses short labels instead of IP lookups to forward packets through the network. |
| Multi-Homed Provider Edge (MH-PE) | PE router that connects to a customer site via redundant Ethernet Segments shared with another PE. |
| Provider Edge (PE) | Edge router that connects customer networks (CEs or switches) to the provider MPLS backbone. |
| Provider (P) Router | Core router within the MPLS network that forwards labeled traffic but does not maintain customer VPN routes. |
| Resource Reservation Protocol (RSVP) | MPLS signaling protocol that establishes label-switched paths (LSPs) with traffic engineering (TE) attributes. |
| Root AC | Access circuit configured as Root in an E-Tree service; can send to or receive from any Leaf. |
| Leaf AC | Access circuit configured as Leaf in an E-Tree service; restricted from sending traffic directly to another Leaf. |
| Scenario-2 (SC-2) | Enhanced E-Tree mode in OcNOS where Leaf Label is advertised via BGP to achieve ingress or egress filtering. |

| Key Terms or Acronym | Description |
|-----------------------------------|--|
| Virtual Local Area Network (VLAN) | Logical network segment used to separate broadcast domains within Layer 2 infrastructure. |
| VPN-ID | Identifier that maps a set of interfaces or EVPN instances to a common Layer 2 or Layer 3 VPN service. |
| Route Target-2 (RT-2) | MAC or IP advertisement route carrying E-Tree extended community and Leaf Label. |
| Route Target-1 (RT-1) | Ethernet A-D per ES route carrying Leaf Label information. |
| Split Horizon | Prevents Leaf-to-Leaf forwarding within the same PE. |

Flow-Aware Transport for EVPN over MPLS

Overview

Flow-Aware Transport (FAT) Labels introduce entropy into the MPLS label stack to improve Equal-Cost Multipath (ECMP) forwarding in EVPN networks. By inserting a FAT label as the bottom-of-stack (BOS) below the EVPN service label, per-flow hashing is enabled across core routers, ensuring efficient traffic distribution without requiring transport-layer changes. This label is independent of any modifications in the Transport Label, ensuring end-to-end entropy delivery.



Note: BGP-based FAT label signaling, as defined in RFC 7432, is not supported.

Feature Characteristics

- FAT Label can be configured on Provider Edge (PE) routers for Unicast EVPN traffic, but validation occurs only on the ingress or egress interfaces of Transit (P) nodes. It does not apply to BUM traffic.
- Supports both single-homed and multihomed deployments.
- Applies to EVPN E-LAN, E-LINE, and E-Tree services.
- Maintains compatibility with existing Control Word (CW) support.

Benefits

- **Improved Load Balancing:** Ensures even distribution of traffic across ECMP paths.
- **Higher Throughput:** Maximizes available bandwidth utilization by avoiding congestion hotspots.
- **Lower Latency:** Reduces queuing and congestion delays in the core.
- **Data Center and Large-Scale EVPN Ready:** Addresses east-west traffic patterns in leaf-spine and redundant mesh topologies.
- **No Transport Dependency:** Entropy is provided by FAT label insertion, not by core transport capabilities.

FAT Prerequisites

Before configuring FAT in [Figure 59](#) topology, ensure the following:

- MPLS, LDP, RSVP, and IS-IS (or OSPF) are operational across all PEs and core routers.
- Loopback interfaces for each router are reachable over the IGP.
- Control-word and flow-label options are supported together for unicast traffic.
- EVPN MPLS and Multihoming are globally enabled.
- BGP EVPN sessions are established among all PEs.
- Port-channels or physical links toward the core have MPLS, LDP, RSVP, and IGP enabled.
- [Core Routers \(P1 and P2\)](#) form the MPLS transport fabric between PE nodes.
- [Access Switches \(SW-1 and SW-2\)](#) are configured with VLAN trunks toward PEs.

PE Nodes

[PE1 and PE2 Nodes](#) and [PE3 and PE4 Nodes](#) provide Active-Standby EVPN connectivity and originate the flow-label insertion for FAT-based ECMP or LAG load balancing.

Enable EVPN MPLS and Multihoming

Enables EVPN MPLS operation with active-standby redundancy between PEs.

```
!  
hardware-profile filter evpn-mpls-mh enable  
!  
evpn mpls enable  
evpn mpls multihoming enable  
!  
evpn mpls vtep-ip-global 1.1.1.1  
!  
evpn etree enable scenario-2  
!
```

Configure Core-Facing Interfaces

Core-facing interfaces carry MPLS, IS-IS, LDP, and RSVP signaling for underlay transport.

```
!  
interface po10  
description connected-to-p1  
load-interval 30  
ip address 12.1.1.1/24  
mtu 9216  
label-switching  
mpls ldp-igp sync ospf holddown-timer 10  
ip router isis 100  
enable-ldp ipv4  
mpls ldp-igp sync-delay 5  
enable-rsvp  
!  
interface po20  
description connected-to-p2  
load-interval 30  
ip address 13.1.1.1/24  
mtu 9216  
label-switching  
mpls ldp-igp sync ospf holddown-timer 10  
ip router isis 100  
enable-ldp ipv4  
mpls ldp-igp sync-delay 5  
enable-rsvp  
!
```

Configure Route Distinguishers and Route Targets

Defines route-target import or export for EVPN services.

```
!  
mac vrf elan_vrf116  
rd 1.1.1.1:116  
route-target both 116:116  
!  
  
!  
mac vrf eline_vrf216
```

```
rd 1.1.1.1:216
route-target both 216:216
!

!
mac vrf etree_vrf316
rd 1.1.1.1:316
route-target both 316:316
!
```

Physical and Port-Channel Interfaces

Prepares physical interfaces and link aggregation for connectivity and redundancy. Configure similar link aggregation in PEs for consistent topology.

- **channel-group:** Aggregates multiple physical interfaces into one logical link for redundancy and higher throughput.
- **mode active:** Configures LACP active mode.
- **speed (g):** Ensures high-speed core links.

```
!
interface ce5
  channel-group 1000 mode active
!
interface xe10
  channel-group 10 mode active
!
interface xe11
  channel-group 10 mode active
!
interface xe12
  speed 10g
  channel-group 20 mode active
!
interface xe26
  channel-group 20 mode active
!
```

Configure LDP and RSVP

Enables MPLS label distribution and RSVP synchronization with entropy label capability for FAT forwarding.

```
!
router ldp
  targeted-peer ipv4 2.2.2.2
  exit-targeted-peer-mode
  targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
  targeted-peer ipv4 4.4.4.4
  exit-targeted-peer-mode
  transport-address ipv4 1.1.1.1
  neighbor 2.2.2.2 auth md5 password plain-text test1
!
router rsvp
!
rsvp-trunk PE1-PE3 ipv4
  to 3.3.3.3
!
rsvp-trunk PE1-PE4 ipv4
  to 4.4.4.4
!
rsvp-trunk PE1-PE2 ipv4
```

```
to 2.2.2.2
!
```

Configure OSPF for MPLS Reachability

OSPF advertises IP routes so MPLS LSPs can be established.

```
!
router ospf 100
  ospf router-id 1.1.1.1
  fast-reroute keep-all-paths
  shutdown
  bfd all-interfaces
  network 1.1.1.1/32 area 0.0.0.0
  network 12.1.1.0/24 area 0.0.0.0
  network 13.1.1.0/24 area 0.0.0.0
!
```

IGP and MPLS Setup

Provides routing and MPLS TE capabilities with FAT support. Configure identical IS-IS or MPLS setup on PEs to enable full TE and FAT ECMP forwarding.

- **mpls traffic-eng router-id:** Identifies router for RSVP and TE.
- **capability cspf:** Enables CSPF for TE path computation.
- **bfd all-interfaces:** Fast failure detection for all interfaces.

```
!
router isis 100
  is-type level-1
  ignore-lsp-errors
  lsp-gen-interval 5
  spf-interval-exp level-1 50 2000
  metric-style wide
  mpls traffic-eng router-id 1.1.1.1
  mpls traffic-eng level-1
  capability cspf
  dynamic-hostname
  bfd all-interfaces
  net 49.0001.0100.0000.1111.00
!
```

Configure BGP for MPLS Services

Supports L3VPN or service connectivity over MPLS or FAT.

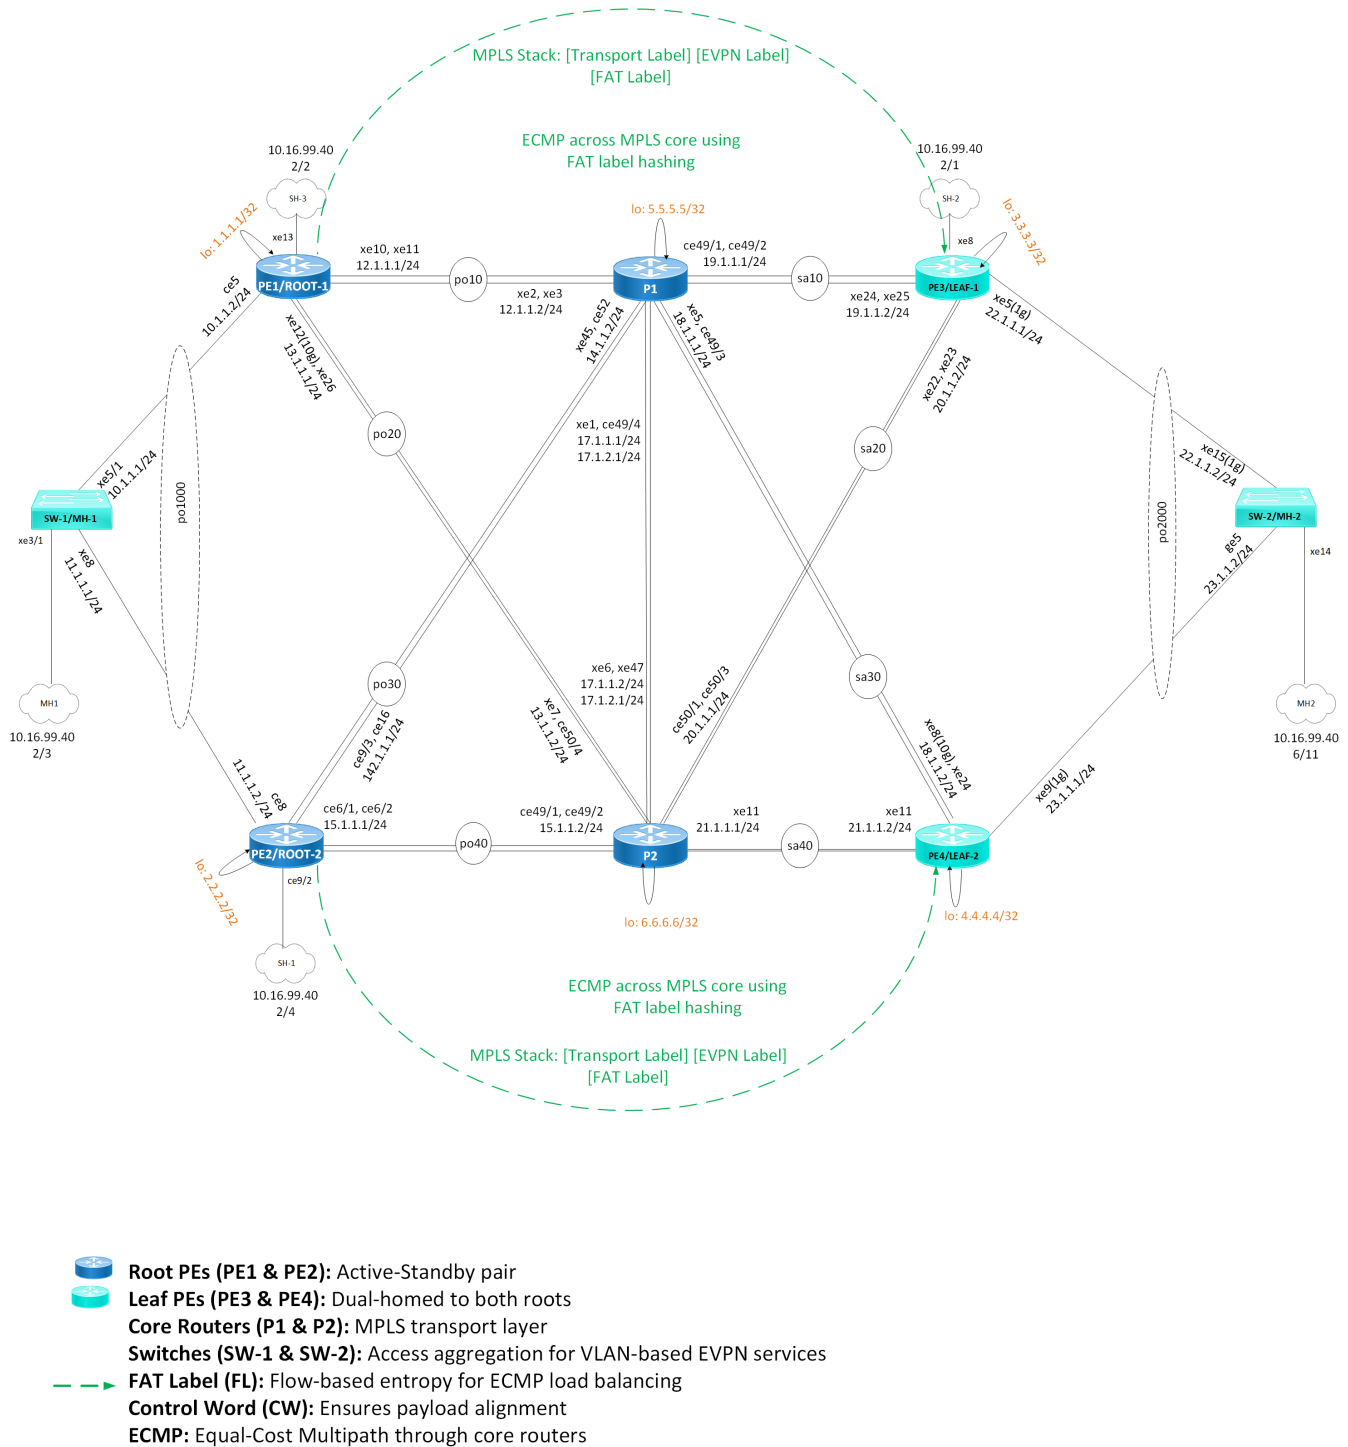
```
!
router bgp 65010
  neighbor PG1 peer-group
  neighbor PG1 remote-as 65010
  neighbor 2.2.2.2 peer-group PG1
  neighbor 3.3.3.3 peer-group PG1
  neighbor 4.4.4.4 peer-group PG1
!
address-family l2vpn evpn
  neighbor PG1 activate
exit-address-family
!
```

FAT Configuration

This section describes a sample use case topology illustrating the FAT label deployment in an EVPN MPLS network. It explains configuring the flow-label on Provider Edge (PE) routers, its operation across the MPLS core, and validation steps to ensure proper functionality and traffic distribution.

Topology

The topology represents an EVPN Active-Standby network where PE1 and PE2 nodes are connect to the PE3 and PE4 nodes through intermediate core routers (P1 and P2) over an MPLS transport network. Each PE is connected to customer edge (CE) switches (SW-1 and SW-2) that provide end-user or service access. Each PE is configured with the `flow-label both` option, which inserts a FAT label beneath the EVPN service label (Bottom of Stack). This FAT label introduces per-flow entropy, allowing Equal-Cost Multipath (ECMP) load balancing across the MPLS core. As a result, traffic between CE switches and the PEs (PE1 and PE2) is efficiently distributed across multiple core paths. The configuration is compatible with control-word support and applies to both EVPN E-LINE and E-LAN services, enhancing bandwidth utilization, reducing congestion, and improving overall network performance.

Figure 59. Flow-Label Enabled EVPN Active-Standby Topology

Note: Before configuration meet all [FAT Prerequisites \(page 1924\)](#).

PE Nodes Configuration



Note: Configure the flow-label option on all participating PE nodes for each EVPN service (E-LINE or E-LAN). This ensures consistent FAT label insertion for a given EVPN instance, enabling end-to-end per-flow entropy and ECMP load balancing across the MPLS core.

EVPN E-LAN and E-LINE Services

Defines EVPN services with optional control-word and flow-label for per-flow ECMP and sequencing. EVPN IDs must match PE's configuration for multihoming.

- **flow-label:** Enables per-flow entropy to improve ECMP load sharing.
- **control-word:** Ensures pseudowire sequencing compatibility.
- **xconnect:** Defines E-LINE pseudowire connection to remote PE.

E-LAN Services

```
!  
evpn mpls id 116 control-word flow-label  
host-reachability-protocol evpn-bgp elan_vrf116  
!
```

E-LINE Services

```
!  
evpn mpls id 5016 xconnect target-mpls-id 6016 control-word flow-label  
host-reachability-protocol evpn-bgp eline_vrf216  
!
```

E-Tree Services

```
!  
evpn mpls id 316 control-word flow-label  
host-reachability-protocol evpn-bgp etree_vrf316  
!
```

Map VLANs to VPN Instances

Associates VLANs or sub-interfaces with EVPN services and multihoming.

- **evpn multi-homed system-mac:** Ensures consistent MAC address for multihomed VLANs.
- **map vpn-id:** Maps access VLAN to the respective EVPN service.

```
!  
interface po1000  
evpn multi-homed system-mac 0000.1212.1313  
!  
  
!  
interface po1000.116 switchport  
description evpn-elan-services-between-pe1-pe2-pe3-pe4-with-CW-and-FL  
encapsulation dot1q 116  
load-interval 30  
access-if-evpn
```

```
    map vpn-id 116
  !

  !
interface po1000.216 switchport
  description evpn-eline-services-MH-MH-with-CW-and-FL
  encapsulation dot1q 216
  load-interval 30
  access-if-evpn
  map vpn-id 5016
  !

  !
interface po1000.316 switchport
  description evpn-etree-services-MH-MH-with-CW-and-FL
  encapsulation dot1q 316
  load-interval 30
  access-if-evpn
  map vpn-id 316
  !
```

FAT Validation

After configuration deployment, the following checks and validations confirm that routing, MPLS, EVPN services, and FAT load-balancing are functioning correctly.

IGP (OSPF or IS-IS) Neighbor Verification

The show commands `show ip ospf neighbor` and `show clns neighbor` confirms that OSPF or IS-IS adjacency is established between PEs and routers. Verifies neighbor state is Full (for OSPF) or Up (for IS-IS), indicating proper routing protocol operation. Ensures loopbacks and links are advertised for MPLS and EVPN tunnels.

MPLS Control Plane Verification (LDP or RSVP)

The show commands `show ldp session` and `show rsvp session` confirms that LDP and RSVP sessions are operational. Ensures MPLS label distribution is complete for all PE-CE and PE-PE links. Verifies RSVP-TE tunnels are ready for traffic engineering and FAT ECMP.

EVPN Control Plane and Tunnel Verification

The following show commands:

- Confirms EVPN session establishment and MPLS tunnels between PEs.
- Checks EVPN service reachability, pseudowire integrity, and MAC address learning.
- Ensures flow-label and control-word settings are active, enabling FAT ECMP load balancing.
- Monitors tunnel counts and traffic statistics to validate correct EVPN distribution.

```
show evpn mpls tunnel summary
show evpn mpls tunnel
show evpn mpls tunnel label
show evpn mpls mac-table
show evpn mpls xconnect
show evpn mpls xconnect tunnel
show evpn mpls xconnect tunnel label
show evpn etree-leaf brief
show bgp l2vpn evpn
show evpn mpls route-count
```

```
show evpn mpls counters <network | evpn-id> <egress | ingress>
```

Post-AC Interface Flap Verification

Steps

- Flap the AC interface (shutdown or no shutdown).
- Verify traffic continuity end-to-end across EVPN services.
- Confirm load-balancing is occurring for services configured with flow-label.

Verification

- Ensures EVPN services maintain redundancy and failover works as expected.
- Confirms flow-label-based per-flow ECMP is distributing traffic evenly.

Packet-Level Verification

Steps

- Capture traffic across EVPN or VPWS tunnels. In this example Wireshark tool capture is used.
- Inspect the MPLS label stack in each packet:
 - Transport Label (Label 0) – Represents IPv4 Explicit Null; ensures the packet is carried with QoS preserved.
 - Service Label (EVPN Label) – Identifies the EVPN ELINE or ELAN instance associated with the payload.
 - Flow Label (FAT Label) – Provides per-flow entropy for ECMP load balancing.
 - Bottom-of-Stack (BoS) bit = 1 – Confirms the last label before payload (Ethernet frame or IP packet).
 - Control-word bit = set – Verifies correct pseudowire encapsulation.

Verification

- Confirms MPLS label stack integrity for each packet.
- Validates label order: Transport → Service → Flow (FAT).
- Ensures ECMP load balancing using the FAT label.
- Confirms the control word and BoS bit are correctly set.
- Verifies that the data plane encapsulation matches the configured EVPN services.

Figure 60. Sample Packet Capture

capture_on_p1_ce49_2_bx_withFL.pcap

File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

Apply a display filter ... <Ctrl-/>

| No. | Time | Source | Destination | Protocol | Length | Info |
|-----|----------|-------------------|-------------------|----------|--------|--------------------------------|
| 1 | 0.000000 | 00:00:00_03:a3:70 | 00:00:00_03:b3:70 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 2 | 0.000169 | 00:00:00_03:a3:c6 | 00:00:00_03:b3:c6 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 3 | 0.008577 | 00:00:00_03:a3:c8 | 00:00:00_03:b3:c8 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 4 | 0.008635 | 00:00:00_03:a3:c9 | 00:00:00_03:b3:c9 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 5 | 0.008798 | 00:00:00_03:a4:1d | 00:00:00_03:b4:1d | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 6 | 0.017434 | 00:00:00_03:a4:6f | 00:00:00_03:b4:6f | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 7 | 0.017542 | 00:00:00_03:a4:71 | 00:00:00_03:b4:71 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 8 | 0.017597 | 00:00:00_03:a4:72 | 00:00:00_03:b4:72 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 9 | 0.026068 | 00:00:00_03:a4:73 | 00:00:00_03:b4:73 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 10 | 0.026125 | 00:00:00_03:a4:7c | 00:00:00_03:b4:7c | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 11 | 0.026181 | 00:00:00_03:a4:7d | 00:00:00_03:b4:7d | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 12 | 0.034784 | 00:00:00_03:a4:ca | 00:00:00_03:b4:ca | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 13 | 0.035012 | 00:00:00_03:a5:20 | 00:00:00_03:b5:20 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 14 | 0.043565 | 00:00:00_03:a5:22 | 00:00:00_03:b5:22 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 15 | 0.043623 | 00:00:00_03:a5:23 | 00:00:00_03:b5:23 | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 16 | 0.043679 | 00:00:00_03:a5:2c | 00:00:00_03:b5:2c | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 17 | 0.043734 | 00:00:00_03:a5:7c | 00:00:00_03:b5:7c | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 18 | 0.043789 | 00:00:00_03:a5:7d | 00:00:00_03:b5:7d | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 19 | 0.052263 | 00:00:00_03:a5:7f | 00:00:00_03:b5:7f | 0x88b5 | 104 | Local Experimental Ethertype 1 |
| 20 | 0.052320 | 00:00:00_03:a5:84 | 00:00:00_03:b5:84 | 0x88b5 | 104 | Local Experimental Ethertype 1 |

> Frame 1: 104 bytes on wire (832 bits), 104 bytes captured (832 bits)

> Ethernet II, Src: EdgecoreNetw_d1:76:79 (b8:6a:97:d1:76:79), Dst: EdgecoreNetw_9f:42:1f (d0:77:ce:9f:42:1f)

> MultiProtocol Label Switching Header, Label: 0 (IPv4 Explicit-Null), Exp: 0, S: 0, TTL: 63

0000 0000 0000 0000 0000 = MPLS Label: IPv4 Explicit-Null (0)

.... = MPLS Experimental Bits: 0

.... = MPLS Bottom Of Label Stack: 0

.... 0011 1111 = MPLS TTL: 63

> MultiProtocol Label Switching Header, Label: 24976, Exp: 0, S: 0, TTL: 64

0000 0110 0001 1001 0000 = MPLS Label: 24976 (0x06190)

.... = MPLS Experimental Bits: 0

.... = MPLS Bottom Of Label Stack: 0

.... 0100 0000 = MPLS TTL: 64

> MultiProtocol Label Switching Header, Label: 576598, Exp: 0, S: 1, TTL: 0

1000 1100 1100 0101 0110 = MPLS Label: 576598 (0x8cc56)

.... = MPLS Experimental Bits: 0

.... = MPLS Bottom Of Label Stack: 1

.... 0000 0000 = MPLS TTL: 0

> Ethernet II, Src: 00:00:00_03:a3:70 (00:00:00:03:a3:70), Dst: 00:00:00_03:b3:70 (00:00:00:03:b3:70)

> Destination: 00:00:00_03:b3:70 (00:00:00:03:b3:70)

> Source: 00:00:00_03:a3:70 (00:00:00:03:a3:70)

Type: 802.1Q Virtual LAN (0x8100)

[Stream index: 1]

> 802.1Q Virtual LAN, PRI: 0, DEI: 0, ID: 115

000. = Priority: Best Effort (default) (0)

...0 = DEI: Ineligible

.... 0000 0111 0011 = ID: 115

Type: Local Experimental Ethertype 1 (0x88b5)

> Data (60 bytes)

Example EVPN Tunnel Status (PE1)

- **Status = Installed:** EVPN tunnel is active and ready.
- **FAT = set:** Flow-label is enabled for per-flow ECMP load balancing.
- Confirms all configured tunnels are active, indicating full connectivity across PEs.

E-LAN

```
PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source           Destination      Status           Up/Down          Update           evpn-id          Local-
Leaf             Remote-Leaf     Ext-Color       FAT
```

```

=====
-----
1.1.1.1      3.3.3.3      Installed    02:27:11    01:16:48    116        ---
   ---      ---      set
.
.
1.1.1.1      2.2.2.2      Installed    02:27:11    01:16:48    116        ---
   ---      ---      set
.
.
1.1.1.1      4.4.4.4      Installed    02:27:11    01:16:48    116        ---
   ---      ---      set

Total number of entries are 75

```

E-LINE

```

PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id Ext-Color FAT
=====
1.1.1.1      2.2.2.2      Installed    03:28:13    01:16:49      5016      6016
   ---      set
1.1.1.1      3.3.3.3      Installed    03:28:13    01:16:49      5016      6016
   ---      set
1.1.1.1      4.4.4.4      Installed    03:28:13    01:16:49      5016      6016
   ---      set
.
.
.

Total number of entries are 12

```

E-Tree

```

PE1#show evpn etree-leaf brief

Leaf
Interface Ifindex  Vnid      Leaf
status  label
-----
po1000.100 0x1f400064 100      Root    ----
po1000.101 0x1f400065 101      Leaf     16
.
.
.
po1000.316 0x1f4000c8 200      Root    ----

Total number of entries are 22

```

FAT Implementation Example

In a data center leaf-spine fabric, multiple equal-cost paths exist between leaf switches (PEs) across the spine (core). Without FAT labels, EVPN traffic uses a single entropy-less service label, limiting load distribution. With FAT labels:

- Each flow receives a unique entropy label.
- Core P-routers hash on the FAT label, distributing traffic across all spine links.
- This achieves near-equal bandwidth utilization and avoids congestion.

EVPN Service Inactive

- Verify FAT label configuration consistency across PEs.
- Ensure control-word settings match between peers.

1. Traffic Not Balanced

- Confirm FAT label is present at BOS in MPLS stack using packet captures.
- Check core router hashing algorithms to ensure they use MPLS labels for entropy.

2. Interoperability Issues

- If the peer does not support FAT labels, disable static FAT configuration to maintain connectivity.

FAT Revised Command

The [evpn mpls id \(page 1936\)](#) command includes the parameter `flow-label` to enable FAT label insertion under EVPN over MPLS. When configured, the router pushes a static FAT label below the EVPN service label, providing per-flow entropy for ECMP load balancing across the MPLS core.

evpn mpls id

Use this command to set the VPN identifier to create an EVPN MPLS tunnel.



Note: To set the VPN identifier for an E-LAN, use the `evpn mpls id` command. For creating an E-LINE/XConnect, use the `evpn mpls xconnect` command with source and target identifiers.

Use the `no` parameter of this command to delete the EVPN MPLS ID for the MPLS tunnel.

Command Syntax

```
evpn mpls id <1-16777215> (| xconnect target-mpls-id <1-16777215>) (|control-word) (|etree-leaf)
(|flow-label)
no evpn mpls id <1-16777215>
```

Parameters

evpn mpls id <1-16777215>

Specifies the EVPN-MPLS tenant identifier. This is a numeric value ranging from 1 to 16777215.

xconnect target-mpls-id <1-16777215>

Enables E-LINE Xconnect. Specifies the target EVID for E-LINE Xconnect. This is a numeric value within the range from 1 to 16777215.

control-word

Enables control-word egress or ingress options for the given EVPN Instance (E-LAN or E-LINE).

etree-leaf

(Optional) Configures the device as a leaf node within the E-Tree topology.

flow-label

Inserts a [Flow-Aware Transport \(FAT\) label](#) for per-flow hashing to improve ECMP traffic distribution.

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 3.0.

Introduced `xconnect target-mpls-id <1-16777215>` parameter in the OcNOS version 4.0, `control-word` parameter in the OcNOS version 6.0.0, `etree-leaf` parameter in the OcNOS version 6.5.1, and `flow-label` parameter in the OcNOS version 7.0.0.

Example

```
(config)#evpn mpls id 10
(config-evpn-mpls)#exit

(config)#no evpn mpls id 10

(config)#evpn mpls id 100 xconnect target-mpls-id 200
(config-evpn-mpls)#exit

(config)#no evpn mpls id 100

(config)#evpn mpls id 300 xconnect target-mpls-id 400 control-word
```

```
(config-evpn-mpls)#exit

(config)#no evpn mpls id 300
```

Use the following command to configure the leaf node as an E-Tree leaf in a MPLS EVPN network.

```
(config)#evpn mpls id 10 etree-leaf
(config-evpn-mpls)#exit
```

The following configuration shows how to enable the `flow-label` parameter under different EVPN MPLS ID instances:

```
!
evpn mpls id 10 xconnect target-mpls-id 15 flow-label
  host-reachability-protocol evpn-bgp vrf2
!
evpn mpls id 20 xconnect target-mpls-id 25 control-word flow-label
  host-reachability-protocol evpn-bgp vrf2
!
evpn mpls id 100 flow-label
  host-reachability-protocol evpn-bgp vrf1
!
evpn mpls id 200 control-word flow-label
  host-reachability-protocol evpn-bgp vrf1
!
```

FAT Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms or Acronym | Description |
|--|---|
| Attachment Circuit (AC) | The physical or logical interface connecting a customer network to a Provider Edge (PE) router. |
| Bidirectional Forwarding Detection (BFD) | A protocol used to detect link failures rapidly between forwarding engines. |
| Bottom of Stack (BoS) | The last label in an MPLS label stack; when set to 1, it identifies the end of the label stack. |
| Control Word (CW) | A 4-byte field added before the payload to preserve packet sequencing and detect packet loss in pseudowires. |
| Equal-Cost Multipath (ECMP) | A routing technique that distributes traffic across multiple equal-cost paths to optimize bandwidth usage. |
| Ethernet LAN (E-LAN) | An EVPN service type that provides multipoint Layer 2 connectivity among multiple PEs. |
| Ethernet Line (E-LINE) | An EVPN service type that provides point-to-point Layer 2 connectivity between two PEs. |
| Ethernet VPN (EVPN) | A BGP-based control-plane mechanism for delivering Layer 2 and Layer 3 VPN services over MPLS or IP networks. |
| Flow-Aware Transport (FAT) | A mechanism using MPLS flow labels to achieve per-flow load balancing across ECMP paths. |
| Flow Label (FL) | A label in the MPLS stack that introduces per-flow entropy for balanced |

| Key Terms or Acronym | Description |
|---------------------------------------|---|
| | distribution across multiple paths. |
| Interior Gateway Protocol (IGP) | A routing protocol (such as OSPF or IS-IS) used within a service provider's network. |
| LLabel Distribution Protocol (LDP) | A protocol that distributes MPLS labels between routers to establish label-switched paths (LSPs). |
| Label Switched Path (LSP) | A unidirectional path established through an MPLS network, identified by labels at each hop. |
| MMultiprotocol Label Switching (MPLS) | A data-carrying technique that forwards packets based on short labels rather than long network addresses. |
| Provider Edge (PE) Router | A router located at the edge of the provider network that connects to customer sites. |
| Provider (P) Router | A core router in the provider network that forwards MPLS traffic between PEs. |
| Pseudowire (PW) | A virtual connection that emulates a point-to-point link over an MPLS core for Layer 2 services. |
| RResource Reservation Protocol (RSVP) | A signaling protocol used to reserve resources and establish label-switched paths in MPLS-TE networks. |

EVPN MAC Limit Enforcement at EVI and AC for EVPN-MPLS

This chapter includes step-by-step configurations for EVPN MAC Limit Enforcement at EVI and AC for EVPN-MPLS.

Overview

The MAC Limit feature in EVPN-MPLS provides a mechanism to control the number of MAC addresses learned at both the Ethernet Virtual Private Network Instance (EVI) level and the Attachment Circuit (AC) level. This feature enhances network efficiency and security by preventing MAC table overflow and mitigating potential denial-of-service attacks in EVPN deployments using MPLS as the underlay, supporting Customer Edge (CE) to Provider Edge (PE) models.

This document outlines the topology, configuration steps, and Command Line Interface (CLI) details for implementing MAC limit enforcement in EVPN-MPLS environments at the EVI and AC levels.

Feature Characteristics

- **MAC Limit Scope:** Limits can be enforced at the EVI level (aggregate MACs across all ACs, BGP-learned MACs, and static MACs) or at the individual AC level (MACs learned from CE to PE).
- **Actions on Limit Breach:** Supports logging and error-disabling for AC-level limits; only logging is supported for EVI-level limits.
- **Threshold Watermarks:** High and low watermark thresholds trigger syslog messages when MAC counts exceed or fall below configured percentages (default high watermark: 90%).
- **Error Disable Recovery:** Configurable recovery timer to automatically re-enable error-disabled ACs after a specified period.
- **Underlay Specific:** The feature is tailored for MPLS underlay in EVPN deployments.
- **Software-Based:** When the MAC limit is reached, further MAC learning is stopped, but flooding continues.

Benefits

- **Prevents MAC Table Overflow:** Improves system stability in EVPN-MPLS environments.
- **Granular Control:** Allows MAC learning limits at both EVI and AC levels.
- **Monitoring Support:** Offers logging for tracking and troubleshooting.
- **Automatic Recovery:** Enables configurable timers to recover from error-disabled states.
- **Enhanced Security:** Limits MAC address learning to improve security in multi-tenant EVPN-MPLS setups.

Prerequisites

- Configure EVPN with MPLS as the overlay protocol.
- Access Interfaces should be configured as switchports with VLAN encapsulation.
- Hardware filters for EVPN-MPLS should be enabled in the multi-homing scenario:

```
hardware-profile filter evpn-mpls-mh enable
```


- Refer to the EVPN-MPLS configuration section in the *Multi-Protocol Label Switching Guide* for EVPN-related configurations.

Limitations

- No platform specific hardware action will be taken on reaching the MAC limit for AC. Only supported action will be error disable.
- With EVI level MAC limit config the only action that would be taken is logging after the limit is reached.

Topology

The EVPN MAC Limit feature can be deployed in any standard EVPN-MPLS topology with CE-PE connectivity.



Configuration

The following configuration steps demonstrate how to enable and configure MAC limit enforcement on PE2 for EVPN-MPLS.

EVPN prerequisite configurations:

PE1:

1. Enable EVPN-MPLS and configure global VTEP IP address.

```
evpn mpls enable
evpn mpls vtep-ip-global 1.1.1.9
```

2. Configure MAC VRF for ELAN service.

```
mac vrf ELAN_602
rd 1.1.1.9:602
route-target both 602:1
evpn mpls id 201
host-reachability-protocol evpn-bgp ELAN_602
```

3. Configure the access interface.

```
interface xe1.201 switchport
encapsulation dot1q 201
load-interval 30
access-if-evpn
map vpn-id 201
```

PE2:

1. Enable EVPN-MPLS and configure global VTEP IP address.

```
evpn mpls enable
evpn mpls vtep-ip-global 1.1.1.40
```


2. Configure MAC VRF for ELAN Service.

```
mac vrf ELAN_602
rd 1.1.1.40:602
route-target both 602:1
evpn mpls id 201
host-reachability-protocol evpn-bgp ELAN_602
```

3. Configure the access interface.

```
interface xe28.201 switchport
encapsulation dot1q 201
load-interval 30
access-if-evpn
map vpn-id 201
```

EVPN-MAC limit profile configuration:

PE2:

Configure a MAC learning limit on the EVPN access interface using a defined profile.

1. Define the MAC Limit Profile.

```
mac-limit-profile p201
learning-limit 10
action log-errdisable
errdisable-timeout 60
```

2. Bind the MAC limit profile to the EVPN access sub-interface.

```
interface xe28.201 switchport
access-if-evpn
learning limit p201
```



Note: Similarly, the MAC limit profile can also be applied at the VNID level.

Configuration snapshot

```
PE2#show running-config interface xe28.201
!
interface xe28.201 switchport
encapsulation dot1q 201
load-interval 30
access-if-evpn
map vpn-id 201
learning limit p201
!
```

Validation

To verify MAC limit enforcement, send traffic with varying numbers of source MAC addresses and observe the system behavior.

- Send Traffic with 10 Source MACs.
- Generate traffic with 10 unique source MACs on xe28.201.
- When the MAC count reaches the high watermark threshold (default 90% of 10 = 9 MACs), a syslog message is generated.

```
PE2#2025 May 31 20:35:00.123 : 7040 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID: 201 Mac limit for AC
xe28.201 High threshold MAC count 9 with high watermark of 9
```

- Send Traffic with 15 Source MACs.
- Generate traffic with 15 unique source MACs, exceeding the learning limit of 10.
- The AC (xe28.201) is error-disabled, and the interface goes down. It recovers after the configured errdisable-timeout (60 seconds).

```
PE2#2025 May 31 20:36:10.456 : 7040 : HSL : CRITI : [EVPN_MAC_LIMIT_2]: VPN-ID: 201 Mac limit for AC
xe28.201 Exceeded MAC count 11 with learning limit of 10
2025 May 31 20:36:10.457 : 7040 : NSM : CRITI : [IFMGR_ERR_DISABLE_DOWN_2]: Attachment Circuit with
the interface xe28.201 on EVPN instance 201 errdisabled successfully due to EVPN-MAC-LIMIT
2025 May 31 20:36:10.458 : 7040 : NSM : CRITI : [IFMGR_IF_DOWN_2]: Interface xe28.201 changed state
to down
```

- Verify Interface Status.

```
PE2#show interface brief
xe28          ETH          --      routed          up      none      10g      --      No  No
xe28.201      SUBINTERFACE --      --          down      ED        10g      --      No  No
```



Notes:

- Action as log-only, is applicable for both AC (attachment circuit) and vnid.
- Action as log-errdisable, is applicable only for AC (attachment circuit).

CLI Commands

This commands section describes the mac limit enforcement at EVI and AC for EVPN-MPLS.

mac-limit-profile

Use this command to create a MAC limit profile to enforce Layer 2 MAC limits. Use the `no` form to delete the profile.



Note: A maximum of 100 mac-limit-profiles can be configured at a time.

Command Syntax

```
mac-limit-profile <PROFILE-NAME>
[no] mac-limit-profile <PROFILE-NAME>
```

Parameters

<PROFILE-NAME>

Name of the MAC limit profile.

Default

None

Command Mode

Configure mode

Applicability

This command is introduced in OcNOS version 6.6.1 version.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#learning-limit 100
OcNOS(config-mac-limit-profile)#high-watermark 80
OcNOS(config-mac-limit-profile)#low-watermark 40
OcNOS(config-mac-limit-profile)#action log-errdisable
OcNOS(config-mac-limit-profile)#errdisable-timeout 60
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#no mac-limit-profile profile_1
OcNOS(config)#commit
```

learning-limit

Use this command to set the learning limit. Use the `no` form to reset to the default limit.

Command Syntax

```
learning-limit <1-131071>
[no] learning-limit
```

Parameters

<1-131071>

Maximum number of MAC addresses (1 to 131071).

Default

131071

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS version 6.6.1 version.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#learning-limit 100
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no learning-limit
OcNOS(config-mac-limit-profile)#commit
```

action

Use this command to set the action type after the MAC limit is reached. If set to log-only, only logs are generated. If set to log-errdisable, both logs and error-disable actions take effect. Use the `no` form to reset to the default action.

Command Syntax

```
action (log-only | log-errdisable)
[no] action
```

Parameters

log-only

Maximum number of MAC addresses (1 to 131071).

log-errdisable

Generates logs and error-disables the interface

Default

log-only

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS version 6.6.1 version.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#action log-errdisable
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no action
OcNOS(config-mac-limit-profile)#commit
```

high-watermark

Use this command to set the high-watermark percentage before the MAC limit is reached. Once the high watermark is reached, an operator log is generated. Use the `no` form to reset to the default.

Command Syntax

```
high-watermark <1-100>
[no] high-watermark
```

Parameters

<1-100>

Percentage of the MAC limit (1 to 100).

Default

90

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS version 6.6.1 version.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#high-watermark 80
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no high-watermark
OcNOS(config-mac-limit-profile)#commit
```

low-watermark

Use this command to set the low watermark percentage. Once the low watermark is reached, an operator log is generated. Use the `no` form to reset to the default.

Command Syntax

```
low-watermark <1-100>
[no] low-watermark
```

Parameters**<1-100>**

Percentage of the MAC limit (1 to 100).

Default

70

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS version 6.6.1 version.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#low-watermark 40
```

```
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no low-watermark
OcNOS(config-mac-limit-profile)#commit
```

errdisable-timeout

Use this command to set the error-disable timeout value in seconds. Once error-disable occurs, this is the duration for which the instance will be operationally shut before being restored. Use the `no` form to reset to the default.

Command Syntax

```
errdisable-timeout <0-86400>
[no] errdisable-timeout
```

Parameters

<0-86400>

Timeout duration in seconds (0 to 86400). A value of 0 disables automatic recovery.

Default

0

Command Mode

MAC-LIMIT-MODE

Applicability

This command is introduced in OcNOS version 6.6.1 version.

Example

```
OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#low-watermark 40
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#mac-limit-profile profile_1
OcNOS(config-mac-limit-profile)#no low-watermark
OcNOS(config-mac-limit-profile)#commit
```

learning limit

Use this command to associate a MAC limit profile with an EVPN-MPLS instance or access interface. Use the `no` form to disassociate the profile.

Command Syntax

```
learning limit <PROFILE-NAME>
[no] learning limit
```

Parameters

<PROFILE-NAME>

Name of the MAC limit profile to apply.

Default

None

Command Mode

EVPN-MPLS-MODE, ACC-IF-EVPN-MODE

Applicability

This command is introduced in OcNOS version 6.6.1 version.

Example

Configure an EVPN-MPLS instance with a mac-limit-profile.

```
OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#mac-limit-profile evi_profile
OcNOS(config-mac-limit-profile)#learning-limit 500
OcNOS(config-mac-limit-profile)#high-watermark 85
OcNOS(config-mac-limit-profile)#low-watermark 45
OcNOS(config-mac-limit-profile)#action log-only
OcNOS(config-mac-limit-profile)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#evpn mpls enable
OcNOS(config)#evpn mpls id 100
OcNOS(config-evpn-mpls)#learning limit evi_profile
OcNOS(config-evpn-mpls)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#evpn mpls id 100
OcNOS(config-evpn-mpls)#no learning limit
OcNOS(config-evpn-mpls)#commit

Configure an EVPN-MPLS access interface with a mac-limit-profile
OcNOS(config)#mac-limit-profile ac_profile
OcNOS(config-mac-limit-profile)#learning-limit 100
OcNOS(config-mac-limit-profile)#high-watermark 92
OcNOS(config-mac-limit-profile)#low-watermark 62
OcNOS(config-mac-limit-profile)#action log-errdisable
OcNOS(config-mac-limit-profile)#errdisable-timeout 120
OcNOS(config-mac-limit-profile)#commit
OcNOS(config-mac-limit-profile)#exit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#interface xel.100 switchport
OcNOS(config-if)#encapsulation dot1q 100
OcNOS(config-if)#access-if-evpn
OcNOS(config-acc-if-evpn)#map vpn-id 100
OcNOS(config-acc-if-evpn)#learning limit ac_profile
OcNOS(config-acc-if-evpn)#commit

OcNOS#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#interface xel.100 switchport
```

```
OcNOS(config-if)#encapsulation dot1q 100
OcNOS(config-if)#access-if-evpn
OcNOS(config-acc-if-evpn)#map vpn-id 100
OcNOS(config-acc-if-evpn)#no learning limit
OcNOS(config-acc-if-evpn)#commit
```

Anycast Gateway Routing for Multiple Subnets in EVPN-IRB

Overview

In the Ethernet VPN Integrated Routing and Bridging (EVPN-IRB) scenario, any two Layer 2 Virtual Network Identifiers (L2 VNID) nodes communicate using the Routing IP Virtual Routing and Forwarding (VRF). This communication is enriched with Anycast Gateway Routing to accommodate communication among multiple subnets under the IRB interface (per VNID).

In the current implementation, the router's primary IPv4 or IPv6 address is either Router Media Access Control (MAC) or Anycast MAC, and the secondary IPv4 or IPv6 address is always the Router MAC address. Hence, Anycast MAC support was only for the primary IP with a single subnet.

Additionally, the BGP router cannot establish a connection with the primary IP as it is in Anycast mode, and the TCP connection is possible only with any of the routers, as both the routes have the IP as Anycast.

To overcome this drawback, the feature is enhanced to configure both Router MAC or Anycast MAC for both primary and secondary subnets.

By default, each subnet uses the Router MAC address received from the ARP/ND cache. The `anycast` argument in `evpn irb-if forwarding anycast gateway` CLI is used to configure the Anycast MAC for primary or secondary subnets. The argument helps to update the ARP/ND cache with Anycast MAC. This enables the user to use Anycast MAC for multiple subnets under L2 VNID. For example, users can have Subnets A, B, C with Anycast MAC and Subnet D with Router MAC.

Feature Characteristics

This feature enhancement provides the following support:

- Enables configuration of either a Router MAC or an Anycast MAC address for primary or secondary subnets.
- Use of Anycast or Routing IP Gateway for multiple subnets under the Layer-2 VNID's.
- Flexibility to have Anycast Gateway for multiple subnets (for example, Subnet A, B, and C) while allowing the other subnet (for example, Subnet D) to be reserved for BGP.
- The InterfaceFull model that provides the flexibility to respond to the ARP/ND requests from the ARP/ND table.
- The InterfaceLess model that use the kernel interface with a unique MAC per interface, either Router MAC or Anycast MAC for all the subnets.

LIMITATIONS:


In InterfaceLess model, the kernel IRB interface has a single MAC that is either Router MAC or Anycast MAC, however, the response message always has Anycast MAC irrespective of whether the interface's IP address is Anycast or Router MAC.

Benefits

Allows users to have primary and secondary subnets with either Router MAC or Anycast MAC. This flexibility provides support for Anycast Gateway for multiple subnets under Layer 2 VNIDs.

Topology

Figure 61. Anycast Gateway support for subnets



Configuration

Following configuration illustrates how to use the `anycast` argument in `evpn irb-if` forwarding `anycast gateway` CLI to configure the Anycast MAC for both primary or secondary subnets.

PE1 Configuration

PE1: Loopback Interface

| | |
|--|--|
| PE1(config)#interface lo | Enter the loopback interface mode. |
| PE1(config-if)#ip address 22.22.22.22/32 secondary | Configure the IP address on loopback interface. |
| PE1(config-if)#ip router isis ISIS-100 | Enable the IS-IS routing on an interface for area 49 (ISIS-100). |
| PE1(config-if)#Commit | Commit the configurations |
| PE1(config-if)#exit | Exit the configuration mode. |

PE1: Global LDP

| | |
|--|--|
| PE1(config)#router ldp | Enter the Router LDP mode. |
| PE1(config-router)#router-id 22.22.22.22 | Enter the LDP router-id. |
| PE1(config-router)#targeted-peer ipv4 29.29.29.29 | Configure the LDP target peer address. |
| PE1(config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted peer mode. |

| | |
|---|--------------------------------------|
| PE1(config-router)#transport-address ipv4 22.22.22.22 | Configure the LDP transport address. |
| PE1(config-router)#Commit | Commit the configurations |
| PE1(config-router)#exit | Exit the configuration mode. |

PE1: Interface Configuration on Network Side

| | |
|--|---|
| PE1(config)#interface po1 | Enter the Interface mode for the port channel interface |
| PE1(config-if)#ip address 10.1.1.22/24 | Configure the IP address on port channel interface |
| PE1(config-if)#label-switching | Enable the label switching |
| PE1(config-if)#ip router isis ISIS-100 | Enable the IS-IS routing on an interface for area 49 (ISIS-100) |
| PE1(config-if)#enable-ldp ipv4 | Enable the LDP IPv4 |
| PE1(config-if)#interface xe22 | Enter interface mode |
| PE1(config-if)#channel-group 1 mode active | Moving interface to Dynamic LAG 1 |
| PE1(config-if)#Commit | Commit the configurations |
| PE1(config-if)#exit | Exit the configuration mode. |

PE1: IGP-ISIS Configuration

| | |
|---|--|
| PE1(config)#router isis ISIS-100 | Create an ISIS routing instance for area 49 (ISIS-100). |
| PE1(config-router)#is-type level-1 | Configure instance as level-1 routing. |
| PE1(config-router)#metric-style wide | Configure the new style of metric type as wide. |
| PE1(config-router)#mpls traffic-eng router-id 22.22.22.22 | Configure MPLS-TE unique router-id TLV. |
| PE1(config-router)#mpls traffic-eng level-1 | Enable the MPLS-TE in is-type Level-1. |
| PE1(config-router)#capability cspf | Enable the Constrained Shortest Path First (CSPF). |
| PE1(config-router)#dynamic-hostname | Configure the host name to be advertised for an ISIS instance. |
| PE1(config-router)#net 49.0001.0000.0000.0001.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| PE1(config-router)#Commit | Commit the configurations |
| PE1(config-router)#exit | Exit the configuration mode. |

PE1: BGP Configuration

| | |
|---|---|
| PE1(config)#router bgp 65535 | Enter into Router BGP mode. |
| PE1(config-router)#bgp router-id 22.22.22.22 | Configure router-id as 22.22.22.22 (loopback ip address). |
| PE1(config-router)#neighbor 29.29.29.29 remote-as 65535 | Configuring PE2 as iBGP neighbor using it's loopback IP. |

| | |
|--|--|
| PE1(config-router)#neighbor 29.29.29.29 update-source lo | Source of routing updates as loopback. |
| PE1(config-router)#neighbor 29.29.29.29 advertisement-interval 0 | Configure advertisement-interval as 0 for fast convergence for PE2 |
| PE1(config-router)#address-family l2vpn evpn | Enter into l2vpn EVPN address family mode. |
| PE1(config-router-af)#neighbor 29.29.29.29 active | Enabling EVPN Address family for neighbor. |
| PE1(config-router-af)#exit-address-family | Exiting of Address family mode. |
| PE1(config-router)#address-family ipv4 vrf ip_vrf205_mgmt | Entering into VRF address family mode. |
| PE1(config-router-af)#redistribute connected | Redistribute connected routes to the network. |
| PE1(config-router-af)#exit-address-family | Exiting of Address family mode. |
| PE1(config-router)#Commit | Commit the configurations |
| PE1(config-router)#exit | Exit the configuration mode. |

PE1: Global EVPN MPLS Command

| | |
|--|---|
| PE1(config)#evpn mpls enable | Enable the EVPN MPLS globally. |
| PE1(config)#evpn mpls irb | Enable the EVPN MPLS IRB globally. |
| PE1(config)#evpn mpls multihoming enable | Enable the Multi homing, save configures and reboot the board for multi homing to be effective. |
| PE1(config)#qos enable | Enable the QOS. |
| PE1(config)#evpn irb-forwarding anycast-gateway-mac 0011.2233.4455 | Configure anycast gateway MAC globally. |
| PE1(config)#evpn mpls vtep-ip-global 22.22.22.22 | Configure VTEP global IP. |
| PE1(config)#Commit | Commit the configurations |
| PE1(config)#exit | Exit the configuration mode. |

PE1: MAC VRF Configuration

| | |
|--|---|
| PE1(config)#mac vrf vrf205_mgmt | Enter Mac VRF mode. |
| PE1(config-vrf)#rd 22.22.22.22:205 | Configuring Route-Distinguisher value. |
| PE1(config-vrf)#route-target both evpn-auto-rt | Configuring import and export value as evpn-auto-rt. Route targets will be derived automatically. |
| PE1(config-vrf)#Commit | Commit the configurations |
| PE1(config-vrf)#exit | Exit the configuration mode. |

PE1: IP VRF Configuration

| | |
|-----------------------------------|-------------------|
| PE1(config)#ip vrf ip_vrf205_mgmt | Enter IP VRF mode |
|-----------------------------------|-------------------|

| | |
|---|--|
| PE1(config-vrf)#rd 22.22.22.22:305 | Configuring Route-Distinguisher value |
| PE1(config-vrf)#route-target both 305:305 | Configuring route target values i.e import and export values |
| PE1(config-vrf)#l3vni 305 | Configure L3 VNID for routing |
| PE1(config-vrf)#Commit | Commit the configurations |
| PE1(config-vrf)#exit | Exit the configuration mode. |

PE1: IRB Interface Configuration with multiple IPs

| | |
|--|---|
| PE1(config)#interface irb127 | Create IRB interface irb127 |
| PE1(config-irb-if)#ip vrf forwarding ip_vrf205_mgmt | Bind the VRF instance to the interface |
| PE1(config-irb-if)#evpn irb-if-forwarding anycast-gateway-mac | Enable an IRB interface to use the global anycast IRB mac- address |
| PE1(config-irb-if)#ip address 98.98.101.1/24 anycast | Configure the IPv4 primary address and use anycast mac address |
| PE1(config-irb-if)#ip address 103.103.102.1/24 secondary | Configure secondary IPv4 secondary address |
| PE1(config-irb-if)#ip address 104.104.103.1/24 secondary anycast | Configure secondary IPv4 secondary address and use as anycast mac address |
| PE1(config-irb-if)#Commit | Commit the configurations |
| PE1(config-irb-if)#exit | Exit the configuration mode. |

PE1: EVPN MPLS Id Configuration

| | |
|---|--|
| PE1(config)#evpn mpls id 127 | Configure secondary IPv4 secondary address |
| PE1(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf205_mgmt | Map the MAC VRF red |
| PE1(config-evpn-mpls)#evpn irb irb127 | Map the IRB interface |
| PE1(config-evpn-mpls)#commit | Commit the configurations |
| PE1(config-evpn-mpls)#exit | Exit the configuration mode. |

PE1: Interface Configuration on Access Side

| | |
|---|--|
| PE1(config)#interface xe72.127 switchport | Creating L2 sub interface of physical interface xe72 |
| PE1(config-if)#encapsulation dot1q 127 | Setting Encapsulation to dot1q with VLAN ID 127 Supported Encapsulation: dot1ad, dot1q, untagged, default |
| PE1(config-if)#rewrite pop | Configure rewrite with action pop |
| PE1(config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| PE1(config-acc-if-evpn)#map vpn-id 127 | Map VPN-ID 127 |

| | |
|--------------------------------|------------------------------|
| PE1(config-acc-if-evpn)#Commit | Commit the configurations |
| PE1(config-acc-if-evpn)#exit | Exit the configuration mode. |

P Configuration

P: Loopback Interface

| | |
|--|--|
| P(config)#interface lo | Enter the Interface mode for the loopback interface. |
| P(config-if)#ip address 8.8.8.8/32 secondary | Configure the IP address on loopback interface. |
| P(config-if)#ip router isis ISIS-100 | Enable the IS-IS routing on an interface for area 49 (ISIS-100). |
| P(config-if)#Commit | Commit the configurations |
| P(config-if)#exit | Exit the configuration mode. |

P: Global LDP

| | |
|---|--------------------------------------|
| P(config)#router ldp | Enter the Router LDP mode. |
| P(config-router)#router-id 8.8.8.8 | Enter the LDP router-id. |
| P(config-router)#transport-address ipv4 8.8.8.8 | Configure the LDP transport address. |
| P(config-router)#Commit | Commit the configurations |
| P(config-router)#exit | Exit the configuration mode. |

P: Interface Configuration on Network Side

| | |
|--|--|
| P(config)#interface pol | Enter the Interface mode for the port channel interface. |
| P(config-if)#ip address 10.1.1.8/24 | Configure the IP address on port channel interface. |
| P(config-if)#label-switching | Enable the label switching. |
| P(config-if)#ip router isis ISIS-100 | Enable the IS-IS routing on an interface for area 49 (ISIS-100). |
| P(config-if)#enable-ldp ipv4 | Enable the LDP IPv4. |
| P(config-if)#Commit | Commit the configurations |
| P(config-if)#exit | Exit the configuration mode. |
| P(config)#interface xe22 | Enter interface mode. |
| P(config-if)#channel-group 1 mode active | Moving interface to Dynamic LAG 1. |
| P(config-if)#Commit | Commit the configurations |
| P(config-if)#exit | Exit the configuration mode. |
| P(config)#interface pol0 | Enter the Interface mode for the port channel interface. |
| P(config-if)#ip address 10.1.3.8/24 | Configure the IP address on port channel interface. |
| P(config-if)#label-switching | Enable the label switching. |
| P(config-if)#ip router isis ISIS-100 | Enable the IS-IS routing on an interface for area 49 |

| | |
|---|-------------------------------------|
| | (ISIS-100). |
| P(config-if)#enable-ldp ipv4 | Enable the LDP IPv4. |
| P(config-if)#interface xe10 | Enter interface mode. |
| P(config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10. |
| P(config-if)#Commit | Commit the configurations |
| P(config-if)#exit | Exit the configuration mode. |

P: IGP-ISIS Configuration

| | |
|---|--|
| P(config)#router isis ISIS-100 | Create an IS-IS routing instance for area 49 (ISIS-100). |
| P(config-router)#is-type level-1 | Configure the instance as level-1 routing. |
| P(config-router)#metric-style wide | Configure the new style of metric type as wide. |
| P(config-router)#mpls traffic-eng router-id 8.8.8.8 | Configure MPLS-TE unique router-id TLV. |
| P(config-router)#mpls traffic-eng level-1 | Enable the MPLS-TE in is-type Level-1. |
| P(config-router)#capability cspf | Enable the CSPF (Constrained Shortest Path First). |
| P(config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| P(config-router)#net 49.0001.0000.0000.0002.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| P(config-router)#Commit | Commit the configurations |
| P(config-router)#exit | Exit the configuration mode. |

PE2 Configuration

PE2: Loopback Interface

| | |
|--|--|
| PE2(config)#interface lo | Enter the Interface mode for the loopback interface. |
| PE2(config-if)#ip address 29.29.29.29/32 secondary | Configure the IP address on loopback interface. |
| PE2(config-if)#ip router isis ISIS-100 | Enable the IS-IS routing on an interface for area 49 (ISIS-100). |
| PE2(config-if)#Commit | Commit the configurations |
| PE2(config-if)#exit | Exit the configuration mode. |

PE2: Global LDP

| | |
|--|--|
| PE2(config)#router ldp | Enter the Router LDP mode. |
| PE2(config-router)#router-id 29.29.29.29 | Enter the LDP router-id. |
| PE2(config-router)#targeted-peer ipv4 22.22.22.22 | Configure the LDP target peer address. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted peer mode. |

| | |
|---|--------------------------------------|
| PE2(config-router)#transport-address ipv4 29.29.29.29 | Configure the LDP transport address. |
| PE2(config-router)#Commit | Commit the configurations |
| PE2(config-router)#exit | Exit the configuration mode. |

PE2: Global LDP

| | |
|--|--|
| PE2(config)#router ldp | Enter the Router LDP mode. |
| PE2(config-router)#router-id 29.29.29.29 | Enter the LDP router-id. |
| PE2(config-router)#targeted-peer ipv4 22.22.22.22 | Configure the LDP target peer address. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted peer mode. |
| PE2(config-router)#transport-address ipv4 29.29.29.29 | Configure the LDP transport address. |
| PE2(config-router)#Commit | Commit the configurations |
| PE2(config-router)#exit | Exit the configuration mode. |

PE2: Interface Configuration on Network Side

| | |
|---|--|
| PE2(config)#interface po10 | Enter the Interface mode for the port channel interface. |
| PE2(config-if)#ip address 10.1.3.29/24 | Configure the IP address on port channel interface. |
| PE2(config-if)#label-switching | Enable the label switching. |
| PE2(config-if)#ip router isis ISIS-100 | Enable the IS-IS routing on an interface for area 49 (ISIS-100). |
| PE2(config-if)#interface gel | Enter interface mode. |
| PE2(config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10. |
| PE2(config-if)#Commit | Commit the configurations |
| PE2(config-if)#exit | Exit the configuration mode. |

PE2: IGP-ISIS Configuration

| | |
|---|---|
| PE2(config)#router isis ISIS-100 | Create an IS-IS routing instance for area 49 (ISIS-100). |
| PE2(config-router)#is-type level-1 | Configure instance as level-1 routing. |
| PE2(config-router)#metric-style wide | Configure the new style of metric type as wide. |
| PE2(config-router)#mpls traffic-eng router-id 29.29.29.29 | Configure MPLS-TE unique router-id TLV. |
| PE2(config-router)#mpls traffic-eng level-1 | Enable the MPLS-TE in is-type Level-1. |
| PE2(config-router)#capability cspf | Enable the CSPF (Constrained Shortest Path First). |
| PE2(config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| PE2(config-router)#net 49.0001.0000.0000.0029.00 | Set a Network Entity Title for this instance, specifying the area |

| | |
|----------------------------|------------------------------|
| | address and the system ID. |
| PE2(config-router) #Commit | Commit the configurations |
| PE2(config-router) #exit | Exit the configuration mode. |

PE2: BGP Configuration

| | |
|---|---|
| PE2(config) #router bgp 65535 | Enter into Router BGP mode. |
| PE2(config-router) #bgp router-id 29.29.29.29 | Configure router-id as 29.29.29.29 (loopback ip address). |
| PE2(config-router) #neighbor 22.22.22.22 remote-as 65535 | Configuring PE2 as iBGP neighbor using it's loopback IP. |
| PE2(config-router) #neighbor 22.22.22.22 update-source lo | Source of routing updates as loopback. |
| PE2(config-router) #neighbor 22.22.22.22 advertisement-interval 0 | Configure advertisement-interval as 0 for fast convergence for PE2. |
| PE2(config-router) #address-family l2vpn evpn | Enter into l2vpn EVPN address family mode. |
| PE2(config-router-af) #neighbor 22.22.22.22 active | Enabling EVPN Address family for neighbor. |
| PE2(config-router-af) #exit-address-family | Exiting of Address family mode. |
| PE2(config-router) #address-family ipv4 vrf ip_vrf205_mgmt | Entering into VRF address family mode. |
| PE2(config-router-af) #redistribute connected | Redistribute connected routes to the network. |
| PE2(config-router-af) #exit-address-family | Exiting of Address family mode. |
| PE2(config-router-af) #commit | Commit the configurations |
| PE2(config-router-af) #exit | Exit the configuration mode. |

PE2: Global EVPN MPLS Command

| | |
|---|--|
| PE2(config) #evpn mpls enable | Enable the EVPN MPLS globally. |
| PE2(config) #evpn mpls irb | Enable the EVPN MPLS IRB globally. |
| PE2(config) #evpn mpls multihoming enable | Enable the Multihoming, save configs and reboot the board for multihoming to be effective. |
| PE2(config) #qos enable | Enable the QOS. |
| PE2(config) #evpn irb-forwarding anycast-gateway-mac 0011.2233.4567 | Configure anycast gateway MAC globally. |
| PE2(config) #evpn mpls vtep-ip-global 29.29.29.29 | Configure VTEP global IP. |
| PE2(config) #Commit | Commit the configurations |
| PE2(config) #exit | Exit the configuration mode. |

PE2: MAC VRF Configuration

| | |
|----------------------------------|---------------------|
| PE2(config) #mac vrf vrf205_mgmt | Enter Mac VRF mode. |
|----------------------------------|---------------------|

| | |
|--|--|
| PE2(config-vrf)#rd 29.29.29.29:205 | Configuring Route-Distinguisher value. |
| PE2(config-vrf)#route-target both evpn-auto-rt | Configuring import and export value as evpn-auto-rt. Route target will be derived automatically. |
| PE2(config-vrf)#Commit | Commit the configurations |
| PE2(config-vrf)#exit | Exit the configuration mode. |

PE2: IP VRF Configuration

| | |
|---|---|
| PE2(config)#ip vrf ip_vrf205_mgmt | Enter IP VRF mode. |
| PE2(config-vrf)#rd 29.29.29.29:305 | Configuring Route-Distinguisher value. |
| PE2(config-vrf)#route-target both 305:305 | Configuring route target values i.e import and export values. |
| PE2(config-vrf)#l3vni 305 | Configure L3 VNID for routing. |
| PE2(config-vrf)#Commit | Commit the configurations |
| PE2(config-vrf)#exit | Exit the configuration mode. |

PE2: IRB Interface Configuration with multiple IPs

| | |
|--|--|
| PE2(config)#interface irb127 | Create IRB interface irb127. |
| PE2(config-irb-if)#ip vrf forwarding ip_vrf205_mgmt | Bind the VRF instance to the interface. |
| PE2(config-irb-if)#evpn irb-if-forwarding anycast-gateway-mac | Enable an IRB interface to use the global anycast IRB mac- address. |
| PE2(config-irb-if)#ip address 99.99.101.1/24 anycast | Configure the IPv4 primary address and use anycast mac address. |
| PE2(config-irb-if)#ip address 103.103.103.1/24 secondary | Configure secondary IPv4 secondary address. |
| PE2(config-irb-if)#ip address 104.104.104.1/24 secondary anycast | Configure secondary IPv4 secondary address and use as anycast mac address. |
| PE2(config-irb-if)#Commit | Commit the configurations |
| PE2(config-irb-if)#exit | Exit the configuration mode. |

PE2: EVPN MPLS Id Configuration

| | |
|---|---|
| PE2(config)#evpn mpls id 127 | Configure secondary IPv4 secondary address. |
| PE2(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf205_mgmt | Map the MAC VRF red. |
| PE2(config-evpn-mpls)#evpn irb irb127 | Map the IRB interface. |
| PE2(config-evpn-mpls)#commit | Commit the configurations |
| PE2(config-evpn-mpls)#exit | Exit the configuration mode. |

PE2: Interface Configuration on Access Side

| | |
|---|---|
| PE2(config)#interface xe12.127 switchport | Creating L2 sub interface on physical interface xe12. |
| PE2(config-if)#encapsulation dot1q 127 | Setting Encapsulation to dot1q with VLAN ID 127 Supported Encapsulation: dot1ad, dot1q, untagged, default. |
| PE2(config-if)#rewrite pop | Configure rewrite with action pop. |
| PE2(config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| PE2(config-acc-if-evpn)#map vpn-id 127 | Map VPN-ID 127. |
| PE2(config-acc-if-evpn)#Commit | Commit the configurations |
| PE2(config-acc-if-evpn)#exit | Exit the configuration mode. |

PE2 Configuration**PE2: Loopback Interface**

| | |
|--|--|
| PE2(config)#interface lo | Enter the Interface mode for the loopback interface. |
| PE2(config-if)#ip address 29.29.29.29/32 secondary | Configure the IP address on loopback interface. |
| PE2(config-if)#ip router isis ISIS-100 | Enable the IS-IS routing on an interface for area 49 (ISIS-100). |
| PE2(config-if)#Commit | Commit the configurations |
| PE2(config-if)#exit | Exit the configuration mode. |

PE2: Global LDP

| | |
|--|--|
| PE2(config)#router ldp | Enter the Router LDP mode. |
| PE2(config-router)#router-id 29.29.29.29 | Enter the LDP router-id. |
| PE2(config-router)#targeted-peer ipv4 22.22.22.22 | Configure the LDP target peer address. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted peer mode. |
| PE2(config-router)#transport-address ipv4 29.29.29.29 | Configure the LDP transport address. |
| PE2(config-router)#Commit | Commit the configurations |
| PE2(config-router)#exit | Exit the configuration mode. |

PE2: Global LDP

| | |
|---|--|
| PE2(config)#router ldp | Enter the Router LDP mode. |
| PE2(config-router)#router-id 29.29.29.29 | Enter the LDP router-id. |
| PE2(config-router)#targeted-peer ipv4 22.22.22.22 | Configure the LDP target peer address. |

| | |
|--|--------------------------------------|
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted peer mode. |
| PE2(config-router)#transport-address ipv4 29.29.29.29 | Configure the LDP transport address. |
| PE2(config-router)#Commit | Commit the configurations |
| PE2(config-router)#exit | Exit the configuration mode. |

PE2: Interface Configuration on Network Side

| | |
|---|--|
| PE2(config)#interface po10 | Enter the Interface mode for the port channel interface. |
| PE2(config-if)#ip address 10.1.3.29/24 | Configure the IP address on port channel interface. |
| PE2(config-if)#label-switching | Enable the label switching. |
| PE2(config-if)#ip router isis ISIS-100 | Enable the IS-IS routing on an interface for area 49 (ISIS-100). |
| PE2(config-if)#enable-ldp ipv4 | Enable the LDP IPv4. |
| PE2(config-if)#interface ge1 | Enter interface mode. |
| PE2(config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10. |
| PE2(config-if)#Commit | Commit the configurations |
| PE2(config-if)#exit | Exit the configuration mode. |

PE2: IGP-ISIS Configuration

| | |
|---|--|
| PE2(config)#router isis ISIS-100 | Create an IS-IS routing instance for area 49 (ISIS-100). |
| PE2(config-router)#is-type level-1 | Configure instance as level-1 routing. |
| PE2(config-router)#metric-style wide | Configure the new style of metric type as wide. |
| PE2(config-router)#mpls traffic-eng router-id 29.29.29.29 | Configure MPLS-TE unique router-id TLV. |
| PE2(config-router)#mpls traffic-eng level-1 | Enable the MPLS-TE in is-type Level-1. |
| PE2(config-router)#capability cspf | Enable the CSPF (Constrained Shortest Path First). |
| PE2(config-router)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance. |
| PE2(config-router)#net 49.0001.0000.0000.0029.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| PE2(config-router)#Commit | Commit the configurations |
| PE2(config-router)#exit | Exit the configuration mode. |

PE2: BGP Configuration

| | |
|--|---|
| PE2(config)#router bgp 65535 | Enter into Router BGP mode. |
| PE2(config-router)#bgp router-id 29.29.29.29 | Configure router-id as 29.29.29.29 (loopback ip address). |

| | |
|--|---|
| PE2(config-router)#neighbor 22.22.22.22 remote-as 65535 | Configuring PE2 as iBGP neighbor using it's loopback IP. |
| PE2(config-router)#neighbor 22.22.22.22 update-source lo | Source of routing updates as loopback. |
| PE2(config-router)#neighbor 22.22.22.22 advertisement-interval 0 | Configure advertisement-interval as 0 for fast convergence for PE2. |
| PE2(config-router)#address-family l2vpn evpn | Enter into l2vpn EVPN address family mode. |
| PE2(config-router-af)#neighbor 22.22.22.22 active | Enabling EVPN Address family for neighbor. |
| PE2(config-router-af)#exit-address-family | Exiting of Address family mode. |
| PE2(config-router)#address-family ipv4 vrf ip_vrf205_mgmt | Entering into VRF address family mode. |
| PE2(config-router-af)#redistribute connected | Redistribute connected routes to the network. |
| PE2(config-router-af)#exit-address-family | Exiting of Address family mode. |
| PE2(config-router-af)#commit | Commit the configurations |
| PE2(config-router-af)#exit | Exit the configuration mode. |

PE2: Global EVPN MPLS Command

| | |
|--|--|
| PE2(config)#evpn mpls enable | Enable the EVPN MPLS globally. |
| PE2(config)#evpn mpls irb | Enable the EVPN MPLS IRB globally. |
| PE2(config)#evpn mpls multihoming enable | Enable the Multihoming, save configs and reboot the board for multihoming to be effective. |
| PE2(config)#qos enable | Enable the QOS. |
| PE2(config)#evpn irb-forwarding anycast-gateway-mac 0011.2233.4567 | Configure anycast gateway MAC globally. |
| PE2(config)#evpn mpls vtep-ip-global 29.29.29.29 | Configure VTEP global IP. |
| PE2(config)#Commit | Commit the configurations |
| PE2(config)#exit | Exit the configuration mode. |

PE2: MAC VRF Configuration

| | |
|--|--|
| PE2(config)#mac vrf vrf205_mgmt | Enter Mac VRF mode. |
| PE2(config-vrf)#rd 29.29.29.29:205 | Configuring Route-Distinguisher value. |
| PE2(config-vrf)#route-target both evpn-auto-rt | Configuring import and export value as evpn-auto-rt. Route target will be derived automatically. |
| PE2(config-vrf)#Commit | Commit the configurations |
| PE2(config-vrf)#exit | Exit the configuration mode. |

PE2: IP VRF Configuration

| | |
|---|---|
| PE2(config)#ip vrf ip_vrf205_mgmt | Enter IP VRF mode. |
| PE2(config-vrf)#rd 29.29.29.29:305 | Configuring Route-Distinguisher value. |
| PE2(config-vrf)#route-target both 305:305 | Configuring route target values i.e import and export values. |
| PE2(config-vrf)#l3vni 305 | Configure L3 VNID for routing. |
| PE2(config-vrf)#Commit | Commit the configurations |
| PE2(config-vrf)#exit | Exit the configuration mode. |

PE2: IRB Interface Configuration with multiple IPs

| | |
|--|--|
| PE2(config)#interface irb127 | Create IRB interface irb127. |
| PE2(config-irb-if)#ip vrf forwarding ip_vrf205_mgmt | Bind the VRF instance to the interface. |
| PE2(config-irb-if)#evpn irb-if-forwarding anycast-gateway-mac | Enable an IRB interface to use the global anycast IRB mac- address. |
| PE2(config-irb-if)#ip address 99.99.101.1/24 anycast | Configure the IPv4 primary address and use anycast mac address. |
| PE2(config-irb-if)#ip address 103.103.103.1/24 secondary | Configure secondary IPv4 secondary address. |
| PE2(config-irb-if)#ip address 104.104.104.1/24 secondary anycast | Configure secondary IPv4 secondary address and use as anycast mac address. |
| PE2(config-irb-if)#Commit | Commit the configurations |
| PE2(config-irb-if)#exit | Exit the configuration mode. |

PE2: EVPN MPLS Id Configuration

| | |
|---|---|
| PE2(config)#evpn mpls id 127 | Configure secondary IPv4 secondary address. |
| PE2(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf205_mgmt | Map the MAC VRF red. |
| PE2(config-evpn-mpls)#evpn irb irb127 | Map the IRB interface. |
| PE2(config-evpn-mpls)#commit | Commit the configurations |
| PE2(config-evpn-mpls)#exit | Exit the configuration mode. |

PE2: Interface Configuration on Access Side

| | |
|---|---|
| PE2(config)#interface xe12.127 switchport | Creating L2 sub interface on physical interface xe12. |
| PE2(config-if)#encapsulation dot1q 127 | Setting Encapsulation to dot1q with VLAN ID 127 Supported Encapsulation: dot1ad, dot1q, untagged, default. |
| PE2(config-if)#rewrite pop | Configure rewrite with action pop. |

| | |
|--|--|
| PE2(config-if)#access-if-evpn | Entering Access mode for EVPN MPLS ID configuration. |
| PE2(config-acc-if-evpn)#map vpn-id 127 | Map VPN-ID 127. |
| PE2(config-acc-if-evpn)#Commit | Commit the configurations |
| PE2(config-acc-if-evpn)#exit | Exit the configuration mode. |

Validation

Verify installed EVPN MPLS tunnels information.

PE1:

```
#show evpn mpls
EVPN-MPLS Information
=====
Codes: NW - Network Port
       AC - Access Port
       (u) - Untagged

VPN-ID   EVI-Name   EVI-Type Type Interface ESI              VLAN   DF-Status Src-
Addr     Dst-Addr
-----
127      ----         L2      NW      ----         ----              ----   ----
      22.22.22.22  29.29.29.29
127      ----         --      AC      xe72.127    --- Single Homed Port ---  ----   ----
      ----
```

PE2:

```
#show evpn mpls tunnel
EVPN-MPLS Network Tunnel Entries
Source           Destination      Status           Up/Down          Update           evpn-id
=====
22.22.22.22      29.29.29.29     Installed        11:40:31         11:40:31         127
```

Verify the MAC addresses that are cached in the EVPN MAC and ARP table.

Verify the Anycast Gateway MAC addresses that are updated when configuring subnets with Anycast MAC:

PE1 verification:

```
#show evpn mpls mac-table
=====
EVPN MPLS MAC Entries
=====
VNID      Interface VlanId   In-VlanId Mac-Addr      VTEP-
Ip/ESI    Type      Status   MAC move AccessPortDesc
-----
127       irb127    ----     ----      0011.2233.4455 22.22.22.22    Static
Local     ----      0        ----
127       irb127    ----     ----      e49d.73b3.c101 22.22.22.22    Static
Local     ----      0        ----
127       ----     ----     ----      0011.2233.4567 29.29.29.29    Static
```

```

Remote ----- 0 -----
127 ----- ---- ---- e8c5.7aff.96de 29.29.29.29 Static
Remote ----- 0 -----
Total number of entries are 4

#show evpn mpls arp-cache
MPLS-EVPN ARP-CACHE Information
=====
EVPN-ID   Ip-Addr           Mac-Addr           Type           Age-Out   Retries-Left
-----
127       98.98.101.1       0011.2233.4455     Static Local    ----
127       99.99.101.1       0011.2233.4567     Static Remote   ----
127       103.103.102.1     e49d.73b3.c101     Static Local    ----
127       103.103.103.1     e8c5.7aff.96de     Static Remote   ----
127       104.104.103.1     0011.2233.4455     Static Local    ----
127       104.104.104.1     0011.2233.4567     Static Remote   ----
Total number of entries are 6

```

PE2 verification:

```

#show evpn mpls mac-table
=====
EVPN MPLS MAC Entries
=====
VNID      Interface VlanId   In-VlanId Mac-Addr   VTEP-      AccessPortDesc
Ip/ESI                               Type        Status      MAC move
-----
127       ----      ----      ----      0011.2233.4455 22.22.22.22 Static
Remote    -----  0         -----
127       irb127    ----      ----      0011.2233.4567 29.29.29.29 Static
Local     -----  0         -----
127       ----      ----      ----      e49d.73b3.c101 22.22.22.22 Static
Remote    -----  0         -----
127       irb127    ----      ----      e8c5.7aff.96de 29.29.29.29 Static
Local     -----  0         -----

Total number of entries are : 4

#show evpn mpls arp-cache
MPLS-EVPN ARP-CACHE Information
=====
ARP Timeout : 180 sec   Random-Jitter-Max : 200

EVPN-ID   Ip-Addr           Mac-Addr           Type           Age-Out   Retries-Left
-----
127       98.98.101.1       0011.2233.4455     Static Remote   ----
127       99.99.101.1       0011.2233.4567     Static Local    ----
127       103.103.102.1     e49d.73b3.c101     Static Remote   ----
127       103.103.103.1     e8c5.7aff.96de     Static Local    ----
127       104.104.103.1     0011.2233.4455     Static Remote   ----
127       104.104.104.1     0011.2233.4567     Static Local    ----

Verify EVPN route count information as per VPN-ID or Route type:

```

Verify EVPN route count information as per VPN-ID or Route type:

PE1 verification:

```

#show evpn mpls route-count
EVPN-MPLS Active route count information
=====
Max supported route count : 131072
Active route count: 8

```


| ----- | ----- | ----- | ----- | ----- |
|-------|-------|---------|---------|---------|
| VNID | Total | MACONLY | MACIPv4 | MACIPv6 |
| ----- | ----- | ----- | ----- | ----- |
| 127 | 8 | 0 | 6 | 2 |

PE2 verification:

```
#show evpn mpls route-count
EVPN-MPLS Active route count information
=====
Max supported route count   : 131072
Active route count: 8
```

| ----- | ----- | ----- | ----- | ----- |
|-------|-------|---------|---------|---------|
| VNID | Total | MACONLY | MACIPv4 | MACIPv6 |
| ----- | ----- | ----- | ----- | ----- |
| 127 | 8 | 0 | 6 | 2 |

Verify in the BGP EVPN table:

PE1 Verification:

```
#show bgp l2vpn evpn
BGP table version is 2, local router ID is 22.22.22.22
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
[EVPN route type]:[ESI]:[VNID]:[relevant route information]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route
```

| Network | Next Hop | Metric | LocPrf | Weight | Path | Peer | Encap |
|---|-------------|--------|--------|--------|------|-------------|-------|
| RD[10:200] | | | | | | | |
| *>i [5]:[0]:[0]:[24]:[80.80.1.0]:[0.0.0.0]:[17] | 29.29.29.29 | 0 | 100 | 0 | ? | 29.29.29.29 | MPLS |
| RD[20:200] | | | | | | | |
| *>i [5]:[0]:[0]:[24]:[80.80.1.0]:[0.0.0.0]:[16] | 10.10.10.10 | 0 | 100 | 0 | ? | 10.10.10.10 | MPLS |
| RD[10.10.10.10:123] | | | | | | | |
| *>i [2]:[0]:[123]:[48,0011:2233:4567]:[32,99.99.99.1]:[22] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [3]:[123]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| RD[10.10.10.10:124] | | | | | | | |
| *>i [2]:[0]:[124]:[48,0011:2233:4567]:[32,99.99.100.1]:[24] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [3]:[124]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| RD[10.10.10.10:125] | | | | | | | |
| *>i [2]:[0]:[125]:[48,0011:2233:4567]:[32,88.88.3.1]:[21] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [3]:[125]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| RD[10.10.10.10:126] | | | | | | | |
| *>i [3]:[126]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |

```

RD[10.10.10.10:205]
*>i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,99.99.101.1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>i  [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>i  [2]:[0]:[127]:[48,d077:ce2a:8001]:[32,103.103.103.1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>i  [2]:[0]:[127]:[48,d077:ce2a:8001]:[128,1100::1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>i  [3]:[127]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS

RD[10.10.10.10:305]
*>i  [5]:[0]:[0]:[24]:[99.99.101.0]:[0.0.0.0]:[17]
      10.10.10.10      0      100      0      ?  10.10.10.10      MPLS
*>i  [5]:[0]:[0]:[24]:[103.103.103.0]:[0.0.0.0]:[17]
      10.10.10.10      0      100      0      ?  10.10.10.10      MPLS
*>i  [5]:[0]:[0]:[24]:[104.104.104.0]:[0.0.0.0]:[17]
      10.10.10.10      0      100      0      ?  10.10.10.10      MPLS

RD[10.10.10.10:333]
*>i  [5]:[0]:[0]:[24]:[99.99.99.0]:[0.0.0.0]:[19]
      10.10.10.10      0      100      0      ?  10.10.10.10      MPLS

RD[10.10.10.10:334]
*>i  [5]:[0]:[0]:[24]:[99.99.100.0]:[0.0.0.0]:[18]
      10.10.10.10      0      100      0      ?  10.10.10.10      MPLS

RD[10.10.10.10:500]
*>i  [3]:[500]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS

RD[22.22.22.22:123] VRF[vrf100_mgmt]:
*>  [2]:[0]:[123]:[48,0011:2233:4455]:[32,98.98.98.1]:[25618]
      22.22.22.22      0      100      32768  i  -----      MPLS
* i  [2]:[0]:[123]:[48,0011:2233:4567]:[32,99.99.99.1]:[28]
      29.29.29.29      0      100      0      i  29.29.29.29      MPLS
* i  10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [3]:[123]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>  [3]:[123]:[32,22.22.22.22]
      22.22.22.22      0      100      32768  i  -----      MPLS
* i  [3]:[123]:[32,29.29.29.29]
      29.29.29.29      0      100      0      i  29.29.29.29      MPLS

RD[22.22.22.22:124] VRF[vrf200_mgmt]:
*>  [2]:[0]:[124]:[48,0011:2233:4455]:[32,98.98.99.1]:[25619]
      22.22.22.22      0      100      32768  i  -----      MPLS
* i  [2]:[0]:[124]:[48,0011:2233:4567]:[32,99.99.100.1]:[29]
      29.29.29.29      0      100      0      i  29.29.29.29      MPLS
* i  10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [3]:[124]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>  [3]:[124]:[32,22.22.22.22]
      22.22.22.22      0      100      32768  i  -----      MPLS
* i  [3]:[124]:[32,29.29.29.29]
      29.29.29.29      0      100      0      i  29.29.29.29      MPLS

RD[22.22.22.22:125] VRF[vrf201_mgmt]:
*>  [2]:[0]:[125]:[48,0011:2233:4455]:[32,88.88.1.1]:[25629]
      22.22.22.22      0      100      32768  i  -----      MPLS
* i  [2]:[0]:[125]:[48,0011:2233:4567]:[32,88.88.3.1]:[37]
      29.29.29.29      0      100      0      i  29.29.29.29      MPLS
* i  10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [3]:[125]:[32,10.10.10.10]

```

```

10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>  [3]:[125]:[32,22.22.22.22]
22.22.22.22      0      100      32768  i  -----      MPLS
* i  [3]:[125]:[32,29.29.29.29]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS

RD[22.22.22.22:126] VRF[vrf202_mgmt]:
* i  [3]:[126]:[32,10.10.10.10]
10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>  [3]:[126]:[32,22.22.22.22]
22.22.22.22      0      100      32768  i  -----      MPLS
* i  [3]:[126]:[32,29.29.29.29]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS

RD[22.22.22.22:200] VRF[blue]:
*>  [2]:[0]:[200]:[48,e49d:73b3:c101]:[32,70.70.1.1]:[25636]
22.22.22.22      0      100      32768  i  -----      MPLS
*>  [3]:[200]:[32,22.22.22.22]
22.22.22.22      0      100      32768  i  -----      MPLS

RD[22.22.22.22:205] VRF[vrf205_mgmt]:
*>  [2]:[0]:[127]:[48,0011:2233:4455]:[32,98.98.101.1]:[25608]
22.22.22.22      0      100      32768  i  -----      MPLS
*>  [2]:[0]:[127]:[48,0011:2233:4455]:[32,104.104.103.1]:[25608]
22.22.22.22      0      100      32768  i  -----      MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,99.99.101.1]:[22]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS
* i  10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[22]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS
* i  10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1][22]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS
* i  10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [2]:[0]:[127]:[48,d077:ce2a:8001]:[32,103.103.103.1]:[26]
10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [2]:[0]:[127]:[48,d077:ce2a:8001]:[128,1100::1][26]
10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>  [2]:[0]:[127]:[48,e49d:73b3:c101]:[32,103.103.102.1]:[25608]
22.22.22.22      0      100      32768  i  -----      MPLS
* i  [2]:[0]:[127]:[48,e8c5:7aff:96de]:[32,103.103.103.1]:[22]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS
* i  [2]:[0]:[127]:[48,e8c5:7aff:96de]:[128,1100::1][22]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS
* i  [3]:[127]:[32,10.10.10.10]
10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>  [3]:[127]:[32,22.22.22.22]
22.22.22.22      0      100      32768  i  -----      MPLS
* i  [3]:[127]:[32,29.29.29.29]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS

RD[22.22.22.22:500] VRF[ELAN_vrf500]:
* i  [3]:[500]:[32,10.10.10.10]
10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>  [3]:[500]:[32,22.22.22.22]
22.22.22.22      0      100      32768  i  -----      MPLS
* i  [3]:[500]:[32,29.29.29.29]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS

RD[22.22.22.22:501] VRF[ELAN_vrf501]:
*>  [3]:[501]:[32,22.22.22.22]
22.22.22.22      0      100      32768  i  -----      MPLS
* i  [3]:[501]:[32,29.29.29.29]
29.29.29.29      0      100      0      i  29.29.29.29      MPLS

RD[22.22.22.22:502] VRF[ELAN_vrf502]:
*>  [3]:[502]:[32,22.22.22.22]
22.22.22.22      0      100      32768  i  -----      MPLS

```

```

* i  [3]:[502]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[22.22.22.22:503] VRF[ELAN_vrf503]:
*>  [3]:[503]:[32,22.22.22.22]
      22.22.22.22          0          100          32768  i  -----    MPLS
* i  [3]:[503]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[22.22.22.22:504] VRF[ELAN_vrf504]:
*>  [3]:[504]:[32,22.22.22.22]
      22.22.22.22          0          100          32768  i  -----    MPLS
* i  [3]:[504]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[22.22.22.22:505] VRF[ELAN_vrf505]:
*>  [3]:[505]:[32,22.22.22.22]
      22.22.22.22          0          100          32768  i  -----    MPLS
* i  [3]:[505]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[22.22.22.22:506] VRF[ELAN_vrf506]:
*>  [3]:[506]:[32,22.22.22.22]
      22.22.22.22          0          100          32768  i  -----    MPLS
* i  [3]:[506]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[22.22.22.22:507] VRF[ELAN_vrf507]:
*>  [3]:[507]:[32,22.22.22.22]
      22.22.22.22          0          100          32768  i  -----    MPLS
* i  [3]:[507]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[22.22.22.22:508] VRF[ELAN_vrf508]:
*>  [3]:[508]:[32,22.22.22.22]
      22.22.22.22          0          100          32768  i  -----    MPLS
* i  [3]:[508]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[22.22.22.22:509] VRF[ELAN_vrf509]:
*>  [3]:[509]:[32,22.22.22.22]
      22.22.22.22          0          100          32768  i  -----    MPLS
* i  [3]:[509]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[22.22.22.22:510] VRF[ELAN_vrf510]:
*>  [3]:[510]:[32,22.22.22.22]
      22.22.22.22          0          100          32768  i  -----    MPLS
* i  [3]:[510]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[22.22.22.22:600] VRF[eline600]:
*>  [1]:[0]:[600]:[25628]
      22.22.22.22          0          100          32768  i  -----    MPLS

RD[22.22.22.22:602] VRF[eline602]:
*>  [1]:[0]:[602]:[25635]
      22.22.22.22          0          100          32768  i  -----    MPLS

RD[22.22.22.22:604] VRF[eline604]:
*>  [1]:[0]:[604]:[25631]
      22.22.22.22          0          100          32768  i  -----    MPLS

RD[22.22.22.22:606] VRF[eline606]:
*>  [1]:[0]:[606]:[25634]
      22.22.22.22          0          100          32768  i  -----    MPLS

RD[22.22.22.22:608] VRF[eline608]:

```

```

*> [1]:[0]:[608]:[25613]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:610] VRF[eline610]:
*> [1]:[0]:[610]:[25626]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:612] VRF[eline612]:
*> [1]:[0]:[612]:[25632]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:614] VRF[eline614]:
*> [1]:[0]:[614]:[25625]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:616] VRF[eline616]:
*> [1]:[0]:[616]:[25612]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:618] VRF[eline618]:
*> [1]:[0]:[618]:[25617]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:620] VRF[eline620]:
*> [1]:[0]:[620]:[25638]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:2002] VRF[vrf2002]:
*> [1]:[0]:[2224]:[25630]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:2003] VRF[vrf2003]:
*> [1]:[0]:[2226]:[25637]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:2004] VRF[vrf2004]:
*> [1]:[0]:[2228]:[25633]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[22.22.22.22:2005] VRF[vrf2005]:
*> [1]:[0]:[2300]:[25639]
      22.22.22.22      0      100      32768 i ----- MPLS

RD[29.29.29.29:123]
*>i [2]:[0]:[123]:[48,0011:2233:4567]:[32,99.99.99.1]:[28]
      29.29.29.29      0      100      0 i 29.29.29.29 MPLS
*>i [3]:[123]:[32,29.29.29.29]
      29.29.29.29      0      100      0 i 29.29.29.29 MPLS

RD[29.29.29.29:124]
*>i [2]:[0]:[124]:[48,0011:2233:4567]:[32,99.99.100.1]:[29]
      29.29.29.29      0      100      0 i 29.29.29.29 MPLS
*>i [3]:[124]:[32,29.29.29.29]
      29.29.29.29      0      100      0 i 29.29.29.29 MPLS

RD[29.29.29.29:125]
*>i [2]:[0]:[125]:[48,0011:2233:4567]:[32,88.88.3.1]:[37]
      29.29.29.29      0      100      0 i 29.29.29.29 MPLS
*>i [3]:[125]:[32,29.29.29.29]
      29.29.29.29      0      100      0 i 29.29.29.29 MPLS

RD[29.29.29.29:126]
*>i [3]:[126]:[32,29.29.29.29]
      29.29.29.29      0      100      0 i 29.29.29.29 MPLS

RD[29.29.29.29:205]
*>i [2]:[0]:[127]:[48,0011:2233:4567]:[32,99.99.101.1]:[22]
      29.29.29.29      0      100      0 i 29.29.29.29 MPLS

```

```

*>i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[22]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS
*>i  [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1]:[22]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS
*>i  [2]:[0]:[127]:[48,e8c5:7aff:96de]:[32,103.103.103.1]:[22]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS
*>i  [2]:[0]:[127]:[48,e8c5:7aff:96de]:[128,1100::1]:[22]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS
*>i  [3]:[127]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:305]
*>i  [5]:[0]:[0]:[24]:[99.99.101.0]:[0.0.0.0]:[18]
      29.29.29.29          0          100          0    ?  29.29.29.29    MPLS
*>i  [5]:[0]:[0]:[24]:[103.103.103.0]:[0.0.0.0]:[18]
      29.29.29.29          0          100          0    ?  29.29.29.29    MPLS
*>i  [5]:[0]:[0]:[24]:[104.104.104.0]:[0.0.0.0]:[18]
      29.29.29.29          0          100          0    ?  29.29.29.29    MPLS

RD[29.29.29.29:333]
*>i  [5]:[0]:[0]:[24]:[99.99.99.0]:[0.0.0.0]:[20]
      29.29.29.29          0          100          0    ?  29.29.29.29    MPLS

RD[29.29.29.29:334]
*>i  [5]:[0]:[0]:[24]:[99.99.100.0]:[0.0.0.0]:[19]
      29.29.29.29          0          100          0    ?  29.29.29.29    MPLS

RD[29.29.29.29:500]
*>i  [3]:[500]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:501]
*>i  [3]:[501]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:502]
*>i  [3]:[502]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:503]
*>i  [3]:[503]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:504]
*>i  [3]:[504]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:505]
*>i  [3]:[505]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:506]
*>i  [3]:[506]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:507]
*>i  [3]:[507]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:508]
*>i  [3]:[508]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:509]
*>i  [3]:[509]:[32,29.29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

RD[29.29.29.29:510]

```

```
*>i  [3]:[510]:[32,29.29.29]
      29.29.29.29          0          100          0    i  29.29.29.29    MPLS

Total number of prefixes 121
```

PE2 Verification:

```
#show bgp l2vpn evpn
BGP table version is 3, local router ID is 29.29.29.29
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevant route information]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route
```

| Network | Next Hop | Metric | LocPrf | Weight | Path | Peer | Encap |
|--|-------------|--------|--------|--------|------|-------------|-------|
| RD[20:200] | | | | | | | |
| *>i [5]:[0]:[0]:[24]:[80.80.1.0]:[0.0.0.0]:[16] | 10.10.10.10 | 0 | 100 | 0 | ? | 10.10.10.10 | MPLS |
| RD[30:200] | | | | | | | |
| *>i [5]:[0]:[0]:[24]:[70.70.1.0]:[0.0.0.0]:[25600] | 22.22.22.22 | 0 | 100 | 0 | ? | 22.22.22.22 | MPLS |
| RD[10.10.10.10:123] | | | | | | | |
| *>i [2]:[0]:[123]:[48,0011:2233:4567]:[32,99.99.99.1]:[22] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [3]:[123]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| RD[10.10.10.10:124] | | | | | | | |
| *>i [2]:[0]:[124]:[48,0011:2233:4567]:[32,99.99.100.1]:[24] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [3]:[124]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| RD[10.10.10.10:125] | | | | | | | |
| *>i [2]:[0]:[125]:[48,0011:2233:4567]:[32,88.88.3.1]:[21] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [3]:[125]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| RD[10.10.10.10:126] | | | | | | | |
| *>i [3]:[126]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| RD[10.10.10.10:200] | | | | | | | |
| *>i [2]:[0]:[100]:[48,d077:ce2a:8001]:[32,80.80.1.1]:[23] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [3]:[100]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| RD[10.10.10.10:205] | | | | | | | |
| *>i [2]:[0]:[127]:[48,0011:2233:4567]:[32,99.99.101.1]:[26] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[26] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1]:[26] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [2]:[0]:[127]:[48,d077:ce2a:8001]:[32,103.103.103.1]:[26] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *>i [2]:[0]:[127]:[48,d077:ce2a:8001]:[128,1100::1]:[26] | | | | | | | |

```

10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>i  [3]:[127]:[32,10.10.10.10]
10.10.10.10      0      100      0      i  10.10.10.10      MPLS

RD[10.10.10.10:305]
*>i  [5]:[0]:[0]:[24]:[99.99.101.0]:[0.0.0.0]:[17]
10.10.10.10      0      100      0      ?  10.10.10.10      MPLS
*>i  [5]:[0]:[0]:[24]:[103.103.103.0]:[0.0.0.0]:[17]
10.10.10.10      0      100      0      ?  10.10.10.10      MPLS
*>i  [5]:[0]:[0]:[24]:[104.104.104.0]:[0.0.0.0]:[17]
10.10.10.10      0      100      0      ?  10.10.10.10      MPLS

RD[10.10.10.10:333]
*>i  [5]:[0]:[0]:[24]:[99.99.99.0]:[0.0.0.0]:[19]
10.10.10.10      0      100      0      ?  10.10.10.10      MPLS

RD[10.10.10.10:334]
*>i  [5]:[0]:[0]:[24]:[99.99.100.0]:[0.0.0.0]:[18]
10.10.10.10      0      100      0      ?  10.10.10.10      MPLS

RD[10.10.10.10:500]
*>i  [3]:[500]:[32,10.10.10.10]
10.10.10.10      0      100      0      i  10.10.10.10      MPLS

RD[22.22.22.22:123]
*>i  [2]:[0]:[123]:[48,0011:2233:4455]:[32,98.98.98.1]:[25618]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>i  [3]:[123]:[32,22.22.22.22]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS

RD[22.22.22.22:124]
*>i  [2]:[0]:[124]:[48,0011:2233:4455]:[32,98.98.99.1]:[25619]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>i  [3]:[124]:[32,22.22.22.22]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS

RD[22.22.22.22:125]
*>i  [2]:[0]:[125]:[48,0011:2233:4455]:[32,88.88.1.1]:[25629]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>i  [3]:[125]:[32,22.22.22.22]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS

RD[22.22.22.22:126]
*>i  [3]:[126]:[32,22.22.22.22]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS

RD[22.22.22.22:205]
*>i  [2]:[0]:[127]:[48,0011:2233:4455]:[32,98.98.101.1]:[25608]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>i  [2]:[0]:[127]:[48,0011:2233:4455]:[32,104.104.103.1]:[25608]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>i  [2]:[0]:[127]:[48,e49d:73b3:c101]:[32,103.103.102.1]:[25608]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>i  [3]:[127]:[32,22.22.22.22]
22.22.22.22      0      100      0      i  22.22.22.22      MPLS

RD[22.22.22.22:305]
*>i  [5]:[0]:[0]:[24]:[98.98.101.0]:[0.0.0.0]:[25601]
22.22.22.22      0      100      0      ?  22.22.22.22      MPLS
*>i  [5]:[0]:[0]:[24]:[103.103.102.0]:[0.0.0.0]:[25601]
22.22.22.22      0      100      0      ?  22.22.22.22      MPLS
*>i  [5]:[0]:[0]:[24]:[104.104.103.0]:[0.0.0.0]:[25601]
22.22.22.22      0      100      0      ?  22.22.22.22      MPLS

RD[22.22.22.22:333]
*>i  [5]:[0]:[0]:[24]:[98.98.98.0]:[0.0.0.0]:[25606]
22.22.22.22      0      100      0      ?  22.22.22.22      MPLS

```



```

RD[22.22.22.22:334]
*>i  [5]:[0]:[0]:[24]:[98.98.99.0]:[0.0.0.0]:[25605]
      22.22.22.22      0      100      0      ?      22.22.22.22      MPLS

RD[22.22.22.22:335]
*>i  [5]:[0]:[0]:[24]:[88.88.1.0]:[0.0.0.0]:[25604]
      22.22.22.22      0      100      0      ?      22.22.22.22      MPLS

RD[22.22.22.22:500]
*>i  [3]:[500]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:501]
*>i  [3]:[501]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:502]
*>i  [3]:[502]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:503]
*>i  [3]:[503]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:504]
*>i  [3]:[504]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:505]
*>i  [3]:[505]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:506]
*>i  [3]:[506]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:507]
*>i  [3]:[507]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:508]
*>i  [3]:[508]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:509]
*>i  [3]:[509]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[22.22.22.22:510]
*>i  [3]:[510]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

RD[29.29.29.29:123] VRF[vrf100_mgmt]:
* i  [2]:[0]:[123]:[48,0011:2233:4455]:[32,98.98.98.1]:[25618]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS
* i  [2]:[0]:[123]:[48,0011:2233:4567]:[32,99.99.99.1]:[22]
      10.10.10.10      0      100      0      i      10.10.10.10      MPLS
*>      29.29.29.29      0      100      32768      i      -----      MPLS
* i  [3]:[123]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i      10.10.10.10      MPLS
* i  [3]:[123]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS
*> [3]:[123]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i      -----      MPLS

RD[29.29.29.29:124] VRF[vrf200_mgmt]:
* i  [2]:[0]:[124]:[48,0011:2233:4455]:[32,98.98.99.1]:[25619]
      22.22.22.22      0      100      0      i      22.22.22.22      MPLS

```

```

* i  [2]:[0]:[124]:[48,0011:2233:4567]:[32,99.99.100.1]:[24]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>    29.29.29.29      0      100      32768  i  -----      MPLS
* i  [3]:[124]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [3]:[124]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>    [3]:[124]:[32,29.29.29.29]
      29.29.29.29      0      100      32768  i  -----      MPLS

RD[29.29.29.29:125] VRF[vrf201_mgmt]:
* i  [2]:[0]:[125]:[48,0011:2233:4455]:[32,88.88.1.1]:[25629]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
* i  [2]:[0]:[125]:[48,0011:2233:4567]:[32,88.88.3.1]:[21]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>    29.29.29.29      0      100      32768  i  -----      MPLS
* i  [3]:[125]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [3]:[125]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>    [3]:[125]:[32,29.29.29.29]
      29.29.29.29      0      100      32768  i  -----      MPLS

RD[29.29.29.29:126] VRF[vrf202_mgmt]:
* i  [3]:[126]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [3]:[126]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>    [3]:[126]:[32,29.29.29.29]
      29.29.29.29      0      100      32768  i  -----      MPLS

RD[29.29.29.29:200] VRF[blue]:
* i  [2]:[0]:[100]:[48,d077:ce2a:8001]:[32,80.80.1.1]:[23]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>    [2]:[0]:[100]:[48,e8c5:7aff:96de]:[32,80.80.1.1]:[38]
      29.29.29.29      0      100      32768  i  -----      MPLS
* i  [3]:[100]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>    [3]:[100]:[32,29.29.29.29]
      29.29.29.29      0      100      32768  i  -----      MPLS

RD[29.29.29.29:205] VRF[vrf205_mgmt]:
* i  [2]:[0]:[127]:[48,0011:2233:4455]:[32,98.98.101.1]:[25608]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4455]:[32,104.104.103.1]:[25608]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,99.99.101.1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>    29.29.29.29      0      100      32768  i  -----      MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>    29.29.29.29      0      100      32768  i  -----      MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
*>    29.29.29.29      0      100      32768  i  -----      MPLS
* i  [2]:[0]:[127]:[48,d077:ce2a:8001]:[32,103.103.103.1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [2]:[0]:[127]:[48,d077:ce2a:8001]:[128,1100::1]:[26]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [2]:[0]:[127]:[48,e49d:73b3:c101]:[32,103.103.102.1]:[25608]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>    [2]:[0]:[127]:[48,e8c5:7aff:96de]:[32,103.103.103.1]:[22]
      29.29.29.29      0      100      32768  i  -----      MPLS
*>    [2]:[0]:[127]:[48,e8c5:7aff:96de]:[128,1100::1]:[22]
      29.29.29.29      0      100      32768  i  -----      MPLS
* i  [3]:[127]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [3]:[127]:[32,22.22.22.22]

```

```

22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[127]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:500] VRF[ELAN_vrf500]:
* i  [3]:[500]:[32,10.10.10.10]
      10.10.10.10      0      100      0      i  10.10.10.10      MPLS
* i  [3]:[500]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[500]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:501] VRF[ELAN_vrf501]:
* i  [3]:[501]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[501]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:502] VRF[ELAN_vrf502]:
* i  [3]:[502]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[502]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:503] VRF[ELAN_vrf503]:
* i  [3]:[503]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[503]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:504] VRF[ELAN_vrf504]:
* i  [3]:[504]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[504]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:505] VRF[ELAN_vrf505]:
* i  [3]:[505]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[505]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:506] VRF[ELAN_vrf506]:
* i  [3]:[506]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[506]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:507] VRF[ELAN_vrf507]:
* i  [3]:[507]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[507]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:508] VRF[ELAN_vrf508]:
* i  [3]:[508]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[508]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:509] VRF[ELAN_vrf509]:
* i  [3]:[509]:[32,22.22.22.22]
      22.22.22.22      0      100      0      i  22.22.22.22      MPLS
*>  [3]:[509]:[32,29.29.29.29]
      29.29.29.29      0      100      32768      i  -----      MPLS

RD[29.29.29.29:510] VRF[ELAN_vrf510]:
* i  [3]:[510]:[32,22.22.22.22]

```

```

                22.22.22.22          0          100          0    i  22.22.22.22    MPLS
*>   [3]:[510]:[32,29.29.29.29]
                29.29.29.29          0          100        32768    i  -----    MPLS

RD[29.29.29.29:650] VRF[vrf650]:
*>   [3]:[650]:[32,29.29.29.29]
                29.29.29.29          0          100        32768    i  -----    MPLS

Total number of prefixes 110

```

Verify the specific type of EVPN routes using VRF:

PE1:

```

#show bgp l2vpn evpn vrf vrf205_mgmt
BGP table version is 1, local router ID is 22.22.22.22
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevant route information]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network          Next Hop          Metric    LocPrf      Weight    Path  Peer          Encap
*>   [2]:[0]:[127]:[48,0011:2233:4455]:[32,98.98.101.1]:[25608]
                22.22.22.22          0          100        32768    i  -----    MPLS
*>   [2]:[0]:[127]:[48,0011:2233:4455]:[32,104.104.103.1]:[25608]
                22.22.22.22          0          100        32768    i  -----    MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,99.99.101.1]:[22]
                29.29.29.29          0          100          0    i  29.29.29.29    MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[22]
                29.29.29.29          0          100          0    i  29.29.29.29    MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[22]
                29.29.29.29          0          100          0    i  29.29.29.29    MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1]:[22]
                29.29.29.29          0          100          0    i  29.29.29.29    MPLS
* i  [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1]:[22]
                29.29.29.29          0          100          0    i  29.29.29.29    MPLS
* i  [2]:[0]:[127]:[48,d077:ce2a:8001]:[32,103.103.103.1]:[26]
                10.10.10.10          0          100          0    i  10.10.10.10    MPLS
* i  [2]:[0]:[127]:[48,d077:ce2a:8001]:[128,1100::1]:[26]
                10.10.10.10          0          100          0    i  10.10.10.10    MPLS
*>   [2]:[0]:[127]:[48,e49d:73b3:c101]:[32,103.103.102.1]:[25608]
                22.22.22.22          0          100        32768    i  -----    MPLS
* i  [2]:[0]:[127]:[48,e8c5:7aff:96de]:[32,103.103.103.1]:[22]
                29.29.29.29          0          100          0    i  29.29.29.29    MPLS
* i  [2]:[0]:[127]:[48,e8c5:7aff:96de]:[128,1100::1]:[22]
                29.29.29.29          0          100          0    i  29.29.29.29    MPLS
* i  [3]:[127]:[32,10.10.10.10]
                10.10.10.10          0          100          0    i  10.10.10.10    MPLS
*>   [3]:[127]:[32,22.22.22.22]
                22.22.22.22          0          100        32768    i  -----    MPLS
* i  [3]:[127]:[32,29.29.29.29]
                29.29.29.29          0          100          0    i  29.29.29.29    MPLS

Total number of prefixes 13

```

PE2:

```

#show bgp l2vpn evpn vrf vrf205_mgmt
BGP table version is 1, local router ID is 29.29.29.29
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,

```

l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevant route information]

- 1 - Ethernet Auto-discovery Route
- 2 - MAC/IP Route
- 3 - Inclusive Multicast Route
- 4 - Ethernet Segment Route
- 5 - Prefix Route

| | Network | Next Hop | Metric | LocPrf | Weight | Path | Peer | Encap |
|-----|--|-------------|--------|--------|--------|------|-------------|-------|
| * i | [2]:[0]:[127]:[48,0011:2233:4455]:[32,98.98.101.1]:[25608] | 22.22.22.22 | 0 | 100 | 0 | i | 22.22.22.22 | MPLS |
| * i | [2]:[0]:[127]:[48,0011:2233:4455]:[32,104.104.103.1]:[25608] | 22.22.22.22 | 0 | 100 | 0 | i | 22.22.22.22 | MPLS |
| * i | [2]:[0]:[127]:[48,0011:2233:4567]:[32,99.99.101.1]:[26] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *> | | 29.29.29.29 | 0 | 100 | 32768 | i | ----- | MPLS |
| * i | [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[26] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *> | | 29.29.29.29 | 0 | 100 | 32768 | i | ----- | MPLS |
| * i | [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1]:[26] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| *> | | 29.29.29.29 | 0 | 100 | 32768 | i | ----- | MPLS |
| * i | [2]:[0]:[127]:[48,d077:ce2a:8001]:[32,103.103.103.1]:[26] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| * i | [2]:[0]:[127]:[48,d077:ce2a:8001]:[128,1100::1]:[26] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| * i | [2]:[0]:[127]:[48,e49d:73b3:c101]:[32,103.103.102.1]:[25608] | 22.22.22.22 | 0 | 100 | 0 | i | 22.22.22.22 | MPLS |
| *> | [2]:[0]:[127]:[48,e8c5:7aff:96de]:[32,103.103.103.1]:[22] | 29.29.29.29 | 0 | 100 | 32768 | i | ----- | MPLS |
| *> | [2]:[0]:[127]:[48,e8c5:7aff:96de]:[128,1100::1]:[22] | 29.29.29.29 | 0 | 100 | 32768 | i | ----- | MPLS |
| * i | [3]:[127]:[32,10.10.10.10] | 10.10.10.10 | 0 | 100 | 0 | i | 10.10.10.10 | MPLS |
| * i | [3]:[127]:[32,22.22.22.22] | 22.22.22.22 | 0 | 100 | 0 | i | 22.22.22.22 | MPLS |
| *> | [3]:[127]:[32,29.29.29.29] | 29.29.29.29 | 0 | 100 | 32768 | i | ----- | MPLS |

Total number of prefixes 13

Verify the specific type of EVPN routes using RD:

PE1:

```
#show bgp l2vpn evpn rd 22.22.22.22:205
BGP table version is 2, local router ID is 22.22.22.22
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

[EVPN route type]:[ESI]:[VNID]:[relevant route information]

- 1 - Ethernet Auto-discovery Route
- 2 - MAC/IP Route
- 3 - Inclusive Multicast Route
- 4 - Ethernet Segment Route
- 5 - Prefix Route

| | Network | Next Hop | Metric | LocPrf | Weight | Path | Peer | Encap |
|---------------------------------------|--|-------------|--------|--------|--------|------|-------|-------|
| RD[22.22.22.22:205] VRF[vrf205_mgmt]: | | | | | | | | |
| *> | [2]:[0]:[127]:[48,0011:2233:4455]:[32,98.98.101.1]:[25608] | 22.22.22.22 | 0 | 100 | 32768 | i | ----- | MPLS |
| *> | [2]:[0]:[127]:[48,0011:2233:4455]:[32,104.104.103.1]:[25608] | 22.22.22.22 | 0 | 100 | 32768 | i | ----- | MPLS |

```

* i [2]:[0]:[127]:[48,0011:2233:4567]:[32,99.99.101.1]:[22]
    29.29.29.29 0 100 0 i 29.29.29.29 MPLS
* i 10.10.10.10 0 100 0 i 10.10.10.10 MPLS
* i [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[22]
    29.29.29.29 0 100 0 i 29.29.29.29 MPLS
* i 10.10.10.10 0 100 0 i 10.10.10.10 MPLS
* i [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1]:[22]
    29.29.29.29 0 100 0 i 29.29.29.29 MPLS
* i 10.10.10.10 0 100 0 i 10.10.10.10 MPLS
* i [2]:[0]:[127]:[48,d077:ce2a:8001]:[32,103.103.103.1]:[26]
    10.10.10.10 0 100 0 i 10.10.10.10 MPLS
* i [2]:[0]:[127]:[48,d077:ce2a:8001]:[128,1100::1]:[26]
    10.10.10.10 0 100 0 i 10.10.10.10 MPLS
*> [2]:[0]:[127]:[48,e49d:73b3:c101]:[32,103.103.102.1]:[25608]
    22.22.22.22 0 100 32768 i ----- MPLS
* i [2]:[0]:[127]:[48,e8c5:7aff:96de]:[32,103.103.103.1]:[22]
    29.29.29.29 0 100 0 i 29.29.29.29 MPLS
* i [2]:[0]:[127]:[48,e8c5:7aff:96de]:[128,1100::1]:[22]
    29.29.29.29 0 100 0 i 29.29.29.29 MPLS
* i [3]:[127]:[32,10.10.10.10]
    10.10.10.10 0 100 0 i 10.10.10.10 MPLS
*> [3]:[127]:[32,22.22.22.22]
    22.22.22.22 0 100 32768 i ----- MPLS
* i [3]:[127]:[32,29.29.29.29]
    29.29.29.29 0 100 0 i 29.29.29.29 MPLS

```

Total number of prefixes 13

PE2:

```

#show bgp l2vpn evpn rd 29.29.29.29:205
BGP table version is 3, local router ID is 29.29.29.29
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevant route information]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

Network          Next Hop          Metric    LocPrf      Weight    Path Peer          Encap
RD[29.29.29.29:205] VRF[vrf205_mgmt]:
* i [2]:[0]:[127]:[48,0011:2233:4455]:[32,98.98.101.1]:[25608]
    22.22.22.22 0 100 0 i 22.22.22.22 MPLS
* i [2]:[0]:[127]:[48,0011:2233:4455]:[32,104.104.103.1]:[25608]
    22.22.22.22 0 100 0 i 22.22.22.22 MPLS
* i [2]:[0]:[127]:[48,0011:2233:4567]:[32,99.99.101.1]:[26]
    10.10.10.10 0 100 0 i 10.10.10.10 MPLS
*> 29.29.29.29 0 100 32768 i ----- MPLS
* i [2]:[0]:[127]:[48,0011:2233:4567]:[32,104.104.104.1]:[26]
    10.10.10.10 0 100 0 i 10.10.10.10 MPLS
*> 29.29.29.29 0 100 32768 i ----- MPLS
* i [2]:[0]:[127]:[48,0011:2233:4567]:[128,1000::1]:[26]
    10.10.10.10 0 100 0 i 10.10.10.10 MPLS
*> 29.29.29.29 0 100 32768 i ----- MPLS
* i [2]:[0]:[127]:[48,d077:ce2a:8001]:[32,103.103.103.1]:[26]
    10.10.10.10 0 100 0 i 10.10.10.10 MPLS
* i [2]:[0]:[127]:[48,d077:ce2a:8001]:[128,1100::1]:[26]
    10.10.10.10 0 100 0 i 10.10.10.10 MPLS
* i [2]:[0]:[127]:[48,e49d:73b3:c101]:[32,103.103.102.1]:[25608]
    22.22.22.22 0 100 0 i 22.22.22.22 MPLS
*> [2]:[0]:[127]:[48,e8c5:7aff:96de]:[32,103.103.103.1]:[22]
    29.29.29.29 0 100 32768 i ----- MPLS
*> [2]:[0]:[127]:[48,e8c5:7aff:96de]:[128,1100::1]:[22]

```

```

                29.29.29.29          0          100          32768 i ----- MPLS
* i  [3]:[127]:[32,10.10.10.10]
                10.10.10.10          0          100          0 i 10.10.10.10 MPLS
* i  [3]:[127]:[32,22.22.22.22]
                22.22.22.22          0          100          0 i 22.22.22.22 MPLS
*>  [3]:[127]:[32,29.29.29.29]
                29.29.29.29          0          100          32768 i ----- MPLS

Total number of prefixes 13

```

Verify the specific type of EVPN routes using Prefix:

PE1:

```

#show bgp l2vpn evpn prefix [3]:[127]:[32,29.29.29.29]
BGP table version is 2, local router ID is 22.22.22.22
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevant route information]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

      Network          Next Hop          Metric    LocPrf      Weight    Path Peer          Encap

RD[22.22.22.22:205] VRF[vrf205_mgmt]:
* i  [3]:[127]:[32,29.29.29.29]
                29.29.29.29          0          100          0 i 29.29.29.29 MPLS

RD[29.29.29.29:205]
*>i  [3]:[127]:[32,29.29.29.29]
                29.29.29.29          0          100          0 i 29.29.29.29 MPLS

Total number of prefixes 2

```

PE2:

```

#show bgp l2vpn evpn prefix [3]:[127]:[32,22.22.22.22]
BGP table version is 3, local router ID is 29.29.29.29
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevant route information]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

      Network          Next Hop          Metric    LocPrf      Weight    Path Peer          Encap

RD[22.22.22.22:205]
*>i  [3]:[127]:[32,22.22.22.22]
                22.22.22.22          0          100          0 i 22.22.22.22 MPLS

RD[29.29.29.29:205] VRF[vrf205_mgmt]:
* i  [3]:[127]:[32,22.22.22.22]
                22.22.22.22          0          100          0 i 22.22.22.22 MPLS

Total number of prefixes 2

```

Verify the specific type of EVPN routes using both VRF and Prefix:

PE1:

```
#show bgp l2vpn evpn vrf vrf205_mgmt prefix [3]:[127]:[32,29.29.29.29]
BGP table version is 1, local router ID is 22.22.22.22
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevant route information]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route
```

| Hop | Network | Metric | LocPrf | Next Weight | Path | Peer | Encap |
|-----|------------------|-------------|--------|-------------|------|------|--------------------|
| * i | [3]:[127]: | | | | | | |
| | [32,29.29.29.29] | 29.29.29.29 | | 0 | 100 | 0 | i 29.29.29.29 MPLS |

Total number of prefixes 1

PE2:

```
#show bgp l2vpn evpn vrf vrf205_mgmt prefix [3]:[127]:[32,22.22.22.22]
BGP table version is 1, local router ID is 29.29.29.29
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevant route information]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route
```

| Hop | Network | Metric | LocPrf | Next Weight | Path | Peer | Encap |
|-----|------------------|-------------|--------|-------------|------|------|--------------------|
| * i | [3]:[127]: | | | | | | |
| | [32,22.22.22.22] | 22.22.22.22 | | 0 | 100 | 0 | i 22.22.22.22 MPLS |

Total number of prefixes 1

Verify detailed information of EVPN routes:

PE1:

```
#show bgp l2vpn evpn vrf vrf205_mgmt prefix [3]:[127]:[32,29.29.29.29] detail

BGP route entry for prefix : [3]:[127]:[32,29.29.29.29]
Route-Distinguisher: [29.29.29.29:205]
Flags : Valid, IBGP
Nexthop : 29.29.29.29 MED value : 0
Community:
Extended Community: RT:65535:1073741951 Encapsulation:MPLS
Weight :0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Mon Oct 9 10:14:47 2023
Peer : 29.29.29.29

Total number of prefixes 1
```


PE2:

```
#show bgp l2vpn evpn vrf vrf205_mgmt prefix [3]:[127]:[32,22.22.22.22] detail

BGP route entry for prefix : [3]:[127]:[32,22.22.22.22]
Route-Distinguisher: [22.22.22.22:205]
Flags : Valid, IBGP
Nexthop : 22.22.22.22 MED value : 0
Community:
Extended Community: RT:65535:1073741951 Encapsulation:MPLS
Weight :0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Mon Apr 15 07:05:47 2019
Peer : 22.22.22.22
```

Abbreviations

Table 86.

| Acronym | Description |
|----------|--|
| VNID | L2 Virtual Network Identifier |
| VRF | Virtual Routing and Forwarding |
| EVPN-IRB | Ethernet VPN Integrated Routing and Bridging |
| MAC | Media Access Control address |

CFM over EVPN-MPLS for ELINE MultiHoming

Overview

The Connectivity Fault Management (CFM) enhances the product offering for the Ethernet LINE (ELINE) services in MultiHoming scenarios. Based on the 802.1ag standard, CFM encompasses Continuity Check Message (CCM), Ping, and Trace functions that help in network fault detection and isolation. This feature extends CFM over EVPN-MPLS from being solely for Single-Homing deployments to a MultiHoming scenario, where a Remote Maintenance End Point (R-MEP) is treated as a single instance by MultiHoming peers.

The [Topology \(page 1984\)](#) illustrates the configuration of User-to-Provider (UP) MEP on PE2 and PE3 Access Circuit (AC) ports, along with the corresponding UP MEP configured on the remote AC port (PE1). This configuration results in the establishment of a CFM session between the PE VTEPs and the remote VTEP.

Feature Characteristics

Functional requirements for CFM over ELINE MultiHoming:

Continuity Check Message

Continuity Check Message (CCM) provides the following capabilities:

- Ensures error-free base configuration for EVPN-MPLS MultiHoming.
- Maintains uniformity of R-MEP and remote-MAC on MultiHoming nodes.
- Enables the data plane to notify the control plane of CCM timeout, port/interface state changes, and Remote Defect Indication (RDI).
- Configures the data plane to send and process CCMs at specified intervals, with options to enable/disable CCM transmission.
- Detects connectivity failures when no CCM frames are received within a set interval and notifies the control plane.
- Programs the data plane to include Port and Interface Status Type-Length-Values (TLVs) in transmitted CCM frames.
- Transmits CC Protocol Data Unit (PDU) frames with IEEE 802.1ag-2007 compliance and supports RDI bit set or reset operations.

Ping and Trace

Ping and Trace provide the following capabilities:

- Facilitates data plane snooping of LBM or Linktrace Message (LTM) received on MEP.
- Traps LTR PDUs received on MEP and processes/replies to LBM received on User-to-Provider (UP) MEP.
- Uplifts CFM PDUs from the data plane to the control plane, and sends CFM PDUs from the control plane to the data plane.
- Provides statistics counters for transmitted Loopback Replies (LBR) and encodes service frame counts in LBM and LBR PDUs.

Benefits

Enhanced Network Monitoring: CFM enables continuous monitoring of network connections, providing real-time insights into connectivity status and performance. This ensures that any issues are quickly detected and addressed.

Quick Fault Detection: Through CCM, the system promptly identifies any disruptions or faults in the network. This swift detection allows for rapid response and minimized downtime.

Efficient Troubleshooting: CFM's Ping and Trace functions help troubleshoot network problems by pinpointing the origin of issues and the paths taken by data packets. This capability streamlines the resolution process.

Robust MultiHoming Support: The extension of CFM support to MultiHoming scenarios ensures that complex network setups remain resilient and well-monitored, even in challenging environments.

Prerequisites

Before configuring and utilizing CFM for ELINE MultiHoming, ensure the following prerequisites are met:

- **Hardware Profiles Configuration**
 - Enable the required hardware profiles to facilitate CFM operations. These include `cfm-domain-name-str`, `cfm-ccm`, and `evpn-mpls-mh` profiles.
 - Establish the hardware-profile filter (`evpn-mpls-mh`) for EVPN-MPLS MultiHoming.
- **EVPN-MPLS Configuration**
 - Enable and configure EVPN MPLS on the relevant devices and enable MultiHoming support within EVPN MPLS.
- **ELINE Service Setup**
 - Establish the ELINE service and assign the corresponding VPN identifiers (VPN-ID).
 - Configure the host-reachability-protocol using EVPN BGP with the associated Virtual Routing and Forwarding (VRF).
- **ELINE AC MultiHoming Configuration**
 - Configure ELINE MultiHoming features with proper encapsulation settings (e.g., `dot1q`) and `access-if-evpn` settings on relevant interfaces.
 - Define the necessary mapping of VPN identifiers (VPN-ID) for the EVPN service.

For more information on the EVPN MPLS configurations, refer to the [EVPN MPLS Commands \(page 1538\)](#) chapters in the Multi-Protocol Label Switching Guide.

- **MAC and MEP Considerations**
 - Ensure that the MEP on MultiHoming nodes has the same MAC. Consistent Media Access Control (MAC) addressing across Access Circuit (AC) ports is essential to facilitate single R-MEP consideration on MultiHoming peers.

For more information on the CFM configurations, refer to the *Carrier Ethernet configuration* section and *CFM and Y.1731 Commands* section in the *OcNOS Carrier Ethernet Guide*.

Meeting these prerequisites ensures a successful setup of CFM for ELINE MultiHoming, enabling enhanced network fault detection and isolation capabilities.


Configuration

This section illustrates the MultiHomed setup for the CFM over EVPN MPLS feature, showcasing examples for ELINE services with LDP as the underlay MPLS path.

Topology

The following topology consists of customer edge routers CE1 and CE2, and with IPv4 Provider Edge routers PE1, PE2, and PE3, all interconnected through the core router P in the IPv4 MPLS provider network.

Figure 62. CFM over EVPN-MPLS for ELINE MH configuration



CFM Configuration

To enhance network management, monitoring, performance, and fault detection, configure the following hardware-profile commands on PE1, PE2, and PE3 devices, and here are the steps for the configurations: [PE1: CFM \(page 1997\)](#) and [PE2/PE3: CFM \(page 1997\)](#).

- Enable the filter for CFM domain name strings with the command `hardware-profile filter cfm-domain-name-str enable`. This filter enhances the network devices ability to process CFM domain name strings, facilitating better network management and service identification.
- Enable statistics collection for CFM Continuity Check Messages (CCM) using the command `hardware-profile statistics cfm-ccm enable`. This feature allows the network devices to gather valuable insights into network performance and fault detection by collecting and analyzing data related to CFM CCMs.

PE1: Loopback Interface

The configuration on PE1 for a loopback interface with IP address 1.1.1.1/32 secondary is set up to provide IP connectivity for the router.

| | |
|--|--|
| PE1#configure terminal | Enter configure mode. |
| PE1(config)#interface lo | Enter the interface mode for the loopback interface lo. |
| PE1(config-if)#ip address 1.1.1.1/32 secondary | Configure a secondary IP address, 1.1.1.1/32, on the loopback interface. |
| PE1(config-if)#exit | Exit interface mode lo. |
| PE1(config)#commit | Commit the transaction. |

PE1: Global LDP

The configuration on PE1 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|--|--|
| PE1(config)#router ldp | Enter the Router LDP mode. |
| PE1(config-router)#router-id 1.1.1.1 | Set the router ID for LDP to 1.1.1.1. |
| PE1(config-router)#targeted-peer ipv4 7.7.7.7 | Configure targeted peer for LDP using IPv4 addresses. |
| PE1(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE1(config-router)#targeted-peer ipv4 8.8.8.8 | Configure targeted peer for LDP using IPv4 addresses. |
| PE1(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE1(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: Global EVPN MPLS Command

The configuration on PE1 for the Global EVPN MPLS includes activating EVPN MPLS defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multihoming, and activating EVPN MPLS multihoming functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|---|--|
| PE1(config)#evpn mpls enable | Activate the EVPN MPLS functionality on PE1, enabling it to participate in EVPN MPLS services. |
| PE1(config)#commit | Commit candidate configuration to be running configuration. |
| PE1(config)#evpn mpls vtep-ip-global 1.1.1.1 | Configure the global VTEP IP address 1.1.1.1, associating it with the loopback IP. |
| PE1(config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter for EVPN MPLS multihoming. |
| PE1(config)#evpn mpls multihoming enable | Activate the EVPN MPLS multihoming functionality, allowing PE1 to support multihomed EVPN MPLS services. |
| PE1(config)#commit | Commit the transaction. |

PE1: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE1 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|---------------------------------------|---|
| PE1(config)#interface xe3 | Enter interface mode xe3. |
| PE1(config-if)#ip address 10.1.0.1/16 | Configure an IP address, 10.1.0.1/16, on the interface xe3. |

| | |
|--------------------------------|--|
| PE1(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE1(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE1(config-if)#exit | Exit interface mode <code>xe3</code> . |
| PE1(config)#commit | Commit the transaction. |

PE1: OSPF Configuration

The below configuration is performed to set up OSPF on PE1, specifying the router ID and defining network interfaces.

| | |
|---|---|
| PE1(config)#router ospf 1 | Enter the router OSPF mode. Configure PE1 to run OSPF with process ID 1. |
| PE1(config-router)#ospf router-id 1.1.1.1 | Set the OSPF router ID to 1.1.1.1, identifying PE1 within the OSPF network. |
| PE1(config-router)#network 1.1.1.1/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| PE1(config-router)#network 10.1.0.0/16 area 0.0.0.0 | Advertise network address in OSPF. |
| PE1(config-router)#exit | Exit router OSPF mode and return to configure mode. |
| PE1(config)#commit | Commit the transaction. |

PE1: BGP Configuration

The below BGP configuration on PE1 is established to enable BGP routing with ASN 1, set the BGP router ID, define iBGP neighbors, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|--|--|
| PE1(config)#router bgp 1 | Enter the Router BGP mode, ASN: 1 |
| PE1(config-router)#bgp router-id 1.1.1.1 | Configure BGP router ID 1.1.1.1, identifying PE1 within the BGP network. |
| PE1(config-router)#neighbor 7.7.7.7 remote-as 1 | Configure neighbor 7.7.7.7 as an iBGP neighbor with their remote AS number 1. |
| PE1(config-router)#neighbor 7.7.7.7 update-source lo | Configure neighbor 7.7.7.7 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE1(config-router)#neighbor 8.8.8.8 remote-as 1 | Configure neighbor 8.8.8.8 as an iBGP neighbor with their remote AS number 1. |
| PE1(config-router)#neighbor 8.8.8.8 update-source lo | Configure neighbor 8.8.8.8 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE1(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE1(config-router-af)#neighbor 7.7.7.7 activate | Activate EVPN for iBGP neighbor 7.7.7.7 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |

| | |
|--|--|
| <code>PE1(config-router-af)#neighbor 8.8.8.8 activate</code> | Activate EVPN for iBGP neighbor 8.8.8.8 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| <code>PE1(config-router-af)#exit</code> | Exit address family mode and return to the router BGP mode. |
| <code>PE1(config-router)#commit</code> | Commit the transaction. |
| <code>PE1(config-router)#exit</code> | Exit router BGP mode and return to the configure mode. |

PE1: MAC VRF Configuration

The below MAC VRF configuration on PE1 is carried out to define and set up VRF named `vrf2` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|--|--|
| <code>PE1(config)#mac vrf vrf2</code> | Enter VRF mode named <code>vrf2</code> . |
| <code>PE1(config-vrf)#rd 1.1.1.1:2</code> | Configure Route-Distinguisher value of 1.1.1.1:2. |
| <code>PE1(config-vrf)#route-target both 2:2</code> | Configure import and export values for the <code>vrf2</code> as 2:2. |
| <code>PE1(config-vrf)#exit</code> | Exit VRF mode and return to the configure mode. |
| <code>PE1(config)#commit</code> | Commit the transaction. |

PE1: EVPN and VRF Mapping

The EVPN and VRF mapping configuration on PE1 establishes mappings between the EVPN identifier and VRF, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|---|---|
| <code>PE1(config)#evpn mpls id 52 xconnect target-mpls-id 2</code> | Configure the EVPN-VPWS identifier with a source identifier of 52 and a target identifier of 2. |
| <code>PE1(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2</code> | Map VRF <code>vrf2</code> to the EVPN-VPWS identifier |
| <code>PE1(config-evpn-mpls)#commit</code> | Commit the transaction. |
| <code>PE1(config-evpn-mpls)#exit</code> | Exit the EVPN MPLS mode and return to the configure mode. |

PE1: Access Port Configuration

The below access port configuration on PE1 is carried out to create a Layer 2 sub-interface within the physical interface, description the interface, configure the encapsulation with VLAN ID, and map VPN-ID for efficient network access and connectivity.

| | |
|---|--|
| <code>PE1(config)#interface xe33</code> | Enter interface mode <code>xe33</code> . |
| <code>PE1(config-if)#interface xe33.2 switchport</code> | Create a Layer 2 sub-interface <code>xe33.2</code> within the physical interface <code>xe33</code> . |
| <code>PE1(config-if)#description access-side-int</code> | Provide a description for the interface. |
| <code>PE1(config-if)#encapsulation dot1q 2</code> | Set encapsulation to dot1q with VLAN ID 2. |

| | |
|---------------------------------------|---|
| PE1(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE1(config-acc-if-evpn)#map vpn-id 52 | Map VPN-ID 52. |
| PE1(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE1(config-if)#exit | Exit interface mode <code>xe33</code> and return to the configure mode. |
| PE1(config)#commit | Commit the transaction. |

P: Loopback Interface

The configuration on P for a loopback interface with IP address `2.2.2.2/32` secondary is set up to provide IP connectivity for the router.

| | |
|--|--|
| P#configure terminal | Enter configure mode. |
| P(config)#interface lo | Enter the interface mode for the loopback interface <code>lo</code> . |
| P(config-if)#ip address 2.2.2.2/32 secondary | Configure a secondary IP address, <code>2.2.2.2/32</code> , on the loopback interface. |
| P(config-if)#exit | Exit interface mode <code>lo</code> . |
| P(config)#commit | Commit the transaction. |

P: Global LDP

The configuration on P for the Global LDP router, specifying router ID to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|------------------------------------|--|
| P(config)#router ldp | Enter the Router LDP mode. |
| P(config-router)#router-id 2.2.2.2 | Set the router ID for LDP to <code>2.2.2.2</code> . |
| P(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| P(config)#commit | Commit the transaction. |

P: Interface Configuration

The below configuration is performed to set up interfaces on P and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|-------------------------------------|--|
| P(config)#interface xe6 | Enter interface mode <code>xe6</code> . |
| P(config-if)#ip address 10.1.0.2/16 | Configure an IP address, <code>10.1.0.2/16</code> , on the interface <code>xe6</code> . |
| P(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P(config-if)#label-switching | Enable label switching on the interface to enable |

| | |
|--------------------------------------|--|
| | MPLS-based packet forwarding. |
| P(config-if)#exit | Exit interface mode xe6. |
| P(config)#commit | Commit the transaction. |
| P(config)#interface xe21 | Enter interface mode xe21. |
| P(config-if)#ip address 123.1.1.1/24 | Configure an IP address, 123.1.1.1/24, on the interface xe21. |
| P(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P(config-if)#exit | Exit interface mode xe21. |
| P(config)#commit | Commit the transaction. |
| P(config)#interface xe20 | Enter interface mode xe20. |
| P(config-if)#ip address 124.1.1.1/24 | Configure an IP address, 124.1.1.1/24, on the interface xe20. |
| P(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| P(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| P(config-if)#exit | Exit interface mode xe20. |
| P(config)#commit | Commit the transaction. |

P: OSPF Configuration

The below configuration is performed to set up OSPF on P, specifying the router ID and defining network interfaces for efficient routing.

| | |
|--|---|
| P(config)#router ospf 1 | Enter the router OSPF mode. Configure P to run OSPF with process ID 1. |
| P(config-router)#ospf router-id 2.2.2.2 | Set the OSPF router ID to 2.2.2.2, identifying P within the OSPF network. |
| P(config-router)#network 2.2.2.2/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| P(config-router)#network 10.1.0.2/16 area 0.0.0.0 | Advertise network address in OSPF. |
| P(config-router)#network 123.1.1.1/24 area 0.0.0.0 | Advertise network address in OSPF. |
| P(config-router)#network 124.1.1.1/24 area 0.0.0.0 | Advertise network address in OSPF. |
| P(config-router)#exit | Exit router OSPF mode and return to the configure mode. |
| P(config)#commit | Commit the transaction. |

PE2: Loopback Interface

The configuration on PE2 for a loopback interface with IP address 7.7.7.7/32 secondary is set up to provide IP connectivity for the router.

| | |
|--|--|
| PE2#configure terminal | Enter configure mode. |
| PE2(config)#interface lo | Enter the interface mode for the loopback interface lo. |
| PE2(config-if)#ip address 7.7.7.7/32 secondary | Configure a secondary IP address, 7.7.7.7/32, on the loopback interface. |
| PE2(config-if)#exit | Exit interface mode lo. |
| PE2(config)#commit | Commit the transaction. |

PE2: Global LDP

The configuration on PE2 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|--|--|
| PE2(config)#router ldp | Enter the Router LDP mode. |
| PE2(config-router)#router-id 7.7.7.7 | Set the router ID for LDP to 7.7.7.7. |
| PE2(config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer for LDP using IPv4 addresses. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE2(config-router)#targeted-peer ipv4 8.8.8.8 | Configure targeted peer for LDP using IPv4 addresses. |
| PE2(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE2(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: Global EVPN MPLS Command

The configuration on PE2 for the Global EVPN MPLS, includes activating EVPN MPLS defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multihoming, and activating EVPN MPLS multihoming functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|---|--|
| PE2(config)#evpn mpls enable | Activate the EVPN MPLS functionality on PE2, enabling it to participate in EVPN MPLS services. |
| PE2(config)#commit | Commit candidate configuration to be running configuration. |
| PE2(config)#evpn mpls vtep-ip-global 7.7.7.7 | Configure the global VTEP IP address 7.7.7.7, associating it with the loopback IP. |
| PE2(config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter for EVPN MPLS multihoming. |

| | |
|--|--|
| PE2(config)#evpn mpls multihoming enable | Activate the EVPN MPLS multihoming functionality, allowing PE2 to support multihomed EVPN MPLS services. |
| PE2(config)#commit | Commit the transaction. |

PE2: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE2 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|--|--|
| PE2(config)#interface xe21 | Enter interface mode xe21. |
| PE2(config-if)#ip address 123.1.1.2/24 | Configure an IP address, 123.1.1.2/24, on the interface xe21. |
| PE2(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE2(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE2(config-if)#exit | Exit interface mode xe21. |
| PE2(config)#commit | Commit the transaction. |

PE2: OSPF Configuration

The below configuration is performed to set up OSPF on PE2, specifying the router ID and defining network interfaces.

| | |
|--|---|
| PE2(config)#router ospf 1 | Enter the router OSPF mode. Configure PE2 to run OSPF with process ID 1. |
| PE2(config-router)#ospf router-id 7.7.7.7 | Set the OSPF router ID to 7.7.7.7, identifying PE2 within the OSPF network. |
| PE2(config-router)#network 7.7.7.7/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| PE2(config-router)#network 123.1.1.0/24 area 0.0.0.0 | Advertise network address in OSPF. |
| PE2(config-router)#exit | Exit router OSPF mode and return to configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: BGP Configuration

The below BGP configuration on PE2 is established to enable BGP routing with ASN 1, set the BGP router ID, define iBGP neighbors, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|---|---|
| PE2(config)#router bgp 1 | Enter the Router BGP mode, ASN: 1 |
| PE2(config-router)#bgp router-id 7.7.7.7 | Configure BGP router ID 7.7.7.7, identifying PE2 within the BGP network. |
| PE2(config-router)#neighbor 1.1.1.1 remote-as 1 | Configure neighbor 1.1.1.1 as an iBGP neighbor with their remote AS number 1. |

| | |
|--|--|
| PE2(config-router)#neighbor 1.1.1.1 update-source lo | Configure neighbor 1.1.1.1 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE2(config-router)#neighbor 8.8.8.8 remote-as 1 | Configure neighbor 8.8.8.8 as an iBGP neighbor with their remote AS number 1. |
| PE2(config-router)#neighbor 8.8.8.8 update-source lo | Configure neighbor 8.8.8.8 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE2(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE2(config-router-af)#neighbor 1.1.1.1 activate | Activate EVPN for iBGP neighbor 1.1.1.1 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE2(config-router-af)#neighbor 8.8.8.8 activate | Activate EVPN for iBGP neighbor 8.8.8.8 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE2(config-router-af)#exit | Exit address family mode and return to the router BGP mode. |
| PE2(config-router)#commit | Commit the transaction. |
| PE2(config-router)#exit | Exit router BGP mode and return to the configure mode. |

PE2: MAC VRF Configuration

The below MAC VRF configuration on PE2 is carried out to define and set up VRF named `vrf2` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|---------------------------------------|--|
| PE2(config)#mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE2(config-vrf)#rd 7.7.7.7:2 | Configure Route-Distinguisher value of <code>7.7.7.7:2</code> . |
| PE2(config-vrf)#route-target both 2:2 | Configure import and export values for the <code>vrf2</code> as <code>2:2</code> . |
| PE2(config-vrf)#exit | Exit VRF mode and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE2: EVPN and VRF Mapping

The EVPN and VRF mapping configuration on PE2 establishes mappings between the EVPN identifier and VRF, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|--|---|
| PE2(config)#evpn mpls id 2 xconnect target-mpls-id 52 | Configure the EVPN-VPWS identifier with a source identifier of 2 and a target identifier of 52. |
| PE2(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf2 | Map VRF <code>vrf2</code> to the EVPN-VPWS identifier |
| PE2(config-evpn-mpls)#commit | Commit the transaction. |
| PE2(config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE2: Access Port Configuration

The below access port configuration on PE2 is carried out to create a Layer 2 sub-interface within the port channel interface, set the load balancing, configure system MAC and the encapsulation with VLAN ID, map VPN-ID for efficient network access and connectivity.

| | |
|---|---|
| PE2(config)#interface po1 | Enter the port channel interface mode for po1 |
| PE2(config-if)#load-interval 30 | Set the load interval to 30. |
| PE2(config-if)#evpn multi-homed system-mac 0000.aaaa.bbbc | Configure the system-mac address 0000.aaaa.bbbc which plays a role in load balancing. |
| PE2(config-if)#interface po1.2 switchport | Create a Layer 2 sub-interface po1.2 within the port channel. |
| PE2(config-if)#encapsulation dot1q 2 | Set encapsulation to dot1q with VLAN ID 2. |
| PE2(config-if)#access-if-evpn | Enter the access mode for EVPN MPLS ID configuration. |
| PE2(config-acc-if-evpn)#map vpn-id 2 | Map VPN-ID 2. |
| PE2(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE2(config-if)#exit | Exit interface mode po1 and return to the configure mode. |
| PE2(config)#commit | Commit the transaction. |

PE3: Loopback Interface

The configuration on PE3 for a loopback interface with IP address 8.8.8.8/32 secondary is set up to provide IP connectivity for the router.

| | |
|--|--|
| PE3#configure terminal | Enter configure mode. |
| PE3(config)#interface lo | Enter the interface mode for the loopback interface lo. |
| PE3(config-if)#ip address 8.8.8.8/32 secondary | Configure a secondary IP address, 8.8.8.8/32, on the loopback interface. |
| PE3(config-if)#exit | Exit interface mode lo. |
| PE3(config)#commit | Commit the transaction. |

PE3: Global LDP

The configuration on PE3 for the Global LDP router, specifying router ID and targeted peers, is done to set up Label Distribution Protocol (LDP) settings for MPLS.

| | |
|---|---|
| PE3(config)#router ldp | Enter the Router LDP mode. |
| PE3(config-router)#router-id 8.8.8.8 | Set the router ID for LDP to 7.7.7.7. |
| PE3(config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer for LDP using IPv4 addresses. |

| | |
|--|--|
| PE3(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE3(config-router)#targeted-peer ipv4 7.7.7.7 | Configure targeted peer for LDP using IPv4 addresses. |
| PE3(config-router-targeted-peer)#exit-targeted-peer-mode | Exit router targeted-peer-mode. |
| PE3(config-router)#exit | Exit router LDP mode and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: Global EVPN MPLS Command

The configuration on PE3 for the Global EVPN MPLS, includes activating EVPN MPLS defining the global VTEP IP address, enabling hardware profile filtering for EVPN MPLS multihoming, and activating EVPN MPLS multihoming functionality, all of which are crucial for enabling EVPN MPLS features.

| | |
|---|--|
| PE3(config)#evpn mpls enable | Activate the EVPN MPLS functionality on PE3, enabling it to participate in EVPN MPLS services. |
| PE3(config)#commit | Commit candidate configuration to be running configuration. |
| PE3(config)#evpn mpls vtep-ip-global 8.8.8.8 | Configure the global VTEP IP address 8.8.8.8, associating it with the loopback IP. |
| PE3(config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware-profile filter for EVPN MPLS multihoming. |
| PE3(config)#evpn mpls multihoming enable | Activate the EVPN MPLS multihoming functionality, allowing PE3 to support multihomed EVPN MPLS services. |
| PE3(config)#commit | Commit the transaction. |

PE3: Interface Configuration Network Side

The below configuration is performed to set up network interfaces on PE3 and enable LDP for IPv4, ensuring proper routing and labeling functionality.

| | |
|--|--|
| PE3(config)#interface xe5 | Enter interface mode xe5. |
| PE3(config-if)#ip address 124.1.1.2/24 | Configure an IP address, 124.1.1.2/24, on the interface xe5. |
| PE3(config-if)#enable-ldp ipv4 | Enable LDP on the physical interface, facilitating the exchange of label information between devices in the network. |
| PE3(config-if)#label-switching | Enable label switching on the interface to enable MPLS-based packet forwarding. |
| PE3(config-if)#exit | Exit interface mode xe5. |
| PE3(config)#interface xe4 | Enter the interface mode for xe4. |
| PE3(config-if)#channel-group 1 mode active | Attach LAG interface xe4. |

| | |
|---------------------|--|
| PE3(config-if)#exit | Exit interface mode <code>xe4</code> and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: OSPF Configuration

The below configuration is performed to set up OSPF on PE3, specifying the router ID and defining network interfaces.

| | |
|--|---|
| PE3(config)#router ospf 1 | Enter the router OSPF mode. Configure PE3 to run OSPF with process ID 1. |
| PE3(config-router)#ospf router-id 8.8.8.8 | Set the OSPF router ID to 8.8.8.8, identifying PE3 within the OSPF network. |
| PE3(config-router)#network 8.8.8.8/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| PE3(config-router)#network 124.1.1.0/24 area 0.0.0.0 | Advertise network address in OSPF. |
| PE3(config-router)#exit | Exit router OSPF mode and return to configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE3: BGP Configuration

The below BGP configuration on PE3 is established to enable BGP routing with ASN 1, set the BGP router ID, define iBGP neighbors, and enable the EVPN address family for efficient routing in an EVPN environment.

| | |
|--|--|
| PE3(config)#router bgp 1 | Enter the Router BGP mode, ASN: 1 |
| PE3(config-router)#bgp router-id 8.8.8.8 | Configure BGP router ID 8.8.8.8, identifying PE3 within the BGP network. |
| PE3(config-router)#neighbor 1.1.1.1 remote-as 1 | Configure neighbor 1.1.1.1 as an iBGP neighbor with their remote AS number 1. |
| PE3(config-router)#neighbor 1.1.1.1 update-source lo | Configure neighbor 1.1.1.1 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE3(config-router)#neighbor 7.7.7.7 remote-as 1 | Configure neighbor 7.7.7.7 as an iBGP neighbor with their remote AS number 1. |
| PE3(config-router)#neighbor 7.7.7.7 update-source lo | Configure neighbor 7.7.7.7 as an iBGP neighbor, specifying the source of routing updates as the loopback interface. |
| PE3(config-router)#address-family l2vpn evpn | Enter into address family mode for L2VPN EVPN. |
| PE3(config-router-af)#neighbor 1.1.1.1 activate | Activate EVPN for iBGP neighbor 1.1.1.1 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE3(config-router-af)#neighbor 7.7.7.7 activate | Activate EVPN for iBGP neighbor 7.7.7.7 within the address family mode, ensuring that EVPN address family is enabled for the neighbor. |
| PE3(config-router-af)#exit | Exit address family mode and return to the router |

| | |
|----------------------------|--|
| | BGP mode. |
| PE3(config-router) #commit | Commit the transaction. |
| PE3(config-router) #exit | Exit router BGP mode and return to the configure mode. |

PE3: MAC VRF Configuration

The below MAC VRF configuration on PE3 is carried out to define and set up VRF named `vrf2` with specific Route-Distinguisher (RD) and route-target values, ensuring segregated MAC address spaces for distinct network services.

| | |
|--|--|
| PE3(config) #mac vrf vrf2 | Enter VRF mode named <code>vrf2</code> . |
| PE3(config-vrf) #rd 8.8.8.8:2 | Configure Route-Distinguisher value of 8.8.8.8:2. |
| PE3(config-vrf) #route-target both 2:2 | Configure import and export values for the <code>vrf2</code> as 2:2. |
| PE3(config-vrf) #exit | Exit VRF mode and return to the configure mode. |
| PE3(config) #commit | Commit the transaction. |

PE3: EVPN and VRF Mapping

The EVPN and VRF mapping configuration on PE3 establishes mappings between the EVPN identifier and VRF, facilitating efficient routing and connectivity in an EVPN network environment.

| | |
|---|---|
| PE3(config) #evpn mpls id 2 xconnect target-mpls-id 52 | Configure the EVPN-VPWS identifier with a source identifier of 2 and a target identifier of 52. |
| PE3(config-evpn-mpls) #host-reachability-protocol evpn-bgp vrf2 | Map VRF <code>vrf2</code> to the EVPN-VPWS identifier |
| PE3(config-evpn-mpls) #commit | Commit the transaction. |
| PE3(config-evpn-mpls) #exit | Exit the EVPN MPLS mode and return to the configure mode. |

PE3: Access Port Configuration

The below access port configuration on PE3 is carried out to create a Layer 2 sub-interface within the port channel interface, set the load balancing, configure system MAC and the encapsulation with VLAN ID, map VPN-ID for efficient network access and connectivity.

| | |
|--|---|
| PE3(config) #interface po1 | Enter the port channel interface mode for <code>po1</code> |
| PE3(config-if) #load-interval 30 | Set the load interval to 30. |
| PE3(config-if) #evpn multi-homed system-mac 0000.aaaa.bbbc | Configure the system-mac address 0000.aaaa.bbbc which plays a role in load balancing. |
| PE3(config-if) #interface po1.2 switchport | Create a Layer 2 sub-interface <code>po1.2</code> within the port channel. |
| PE3(config-if) #encapsulation dot1q 2 | Set encapsulation to dot1q with VLAN ID 2. |
| PE3(config-if) #access-if-evpn | Enter the access mode for EVPN MPLS ID |

| | |
|--------------------------------------|--|
| | configuration. |
| PE3(config-acc-if-evpn)#map vpn-id 2 | Map VPN-ID 2. |
| PE3(config-acc-if-evpn)#exit | Exit the access mode and return to the interface mode. |
| PE3(config-if)#exit | Exit interface mode <code>po1</code> and return to the configure mode. |
| PE3(config)#commit | Commit the transaction. |

PE1: CFM

The following configuration enables CFM monitoring and maintenance for EVPN services on PE1.

| | |
|---|---|
| PE1(config)#hardware-profile filter cfm-domain-name-str enable | Enable the CFM domain name as a character string profile for CFM filtering. |
| PE1(config)#ethernet cfm domain-type character-string domain-name evpn1 level 7 mip-creation none | Create a CFM domain for EVPN ELINE with a character string type and set MIP creation to <code>none</code> . |
| PE1(config-ether-cfm)#service ma-type string ma-name evp1 | Define a maintenance association (MA) type with the string <code>evp1</code> . |
| PE1(config-ether-cfm-ma)#evpn 52 | Configure the MA for EVPN ID 52. |
| PE1(config-ether-cfm-ma)#ethernet cfm mep up mpid 20 active true evpn 52 | Create an up-maintenance endpoint (MEP) for local EVPN ID 52. |
| PE1(config-ether-cfm-ma-mep)#cc multicast state enable | Enable multicast continuity check (CC) state for the MEP. |
| PE1(config-ether-cfm-ma-mep)#cc interval 100 | Set the CC interval to 100 milliseconds. |
| PE1(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode | Exit Ethernet MA MEP mode. |
| PE1(config-ether-cfm-ma)#mep crosscheck mpid 10 | Configure cross-check to the remote MEP. |
| PE1(config-ether-cfm-ma)#exit-ether-ma-mode | Exit Ethernet MA mode. |
| PE1(config-ether-cfm)#exit | Exit Ethernet CFM mode and return to the configure mode. |
| PE1(config)#commit | Commit candidate configuration to be running configuration. |

PE2/PE3: CFM

The following configuration enables CFM monitoring and maintenance for EVPN services on PE2 and PE3 devices.



Note: Apply the same set of configurations to the PE3 device.

| | |
|---|---|
| PE2(config)#hardware-profile filter cfm-domain-name-str enable | Enable the CFM domain name as a character string profile for CFM filtering. |
| PE2(config)#ethernet cfm domain-type character-string domain-name evpn1 level 7 mip-creation none | Create a CFM domain for EVPN ELINE with a character string type and set MIP creation to <code>none</code> . |

| | |
|---|--|
| PE2(config-ether-cfm)#service ma-type string ma-name evp1 | Define a maintenance association (MA) type with the string evp1. |
| PE2(config-ether-cfm-ma)#evpn 2 | Configure the MA for EVPN ID 2. |
| PE2(config-ether-cfm-ma)#ethernet cfm mep up mpid 10 active true evpn 2 | Create an up-maintenance endpoint (MEP) for local EVPN ID 2. |
| PE2(config-ether-cfm-ma-mep)#cc multicast state enable | Enable multicast continuity check (CC) state for the MEP. |
| PE2(config-ether-cfm-ma-mep)#cc interval 100 | Set the CC interval to 100 milliseconds. |
| PE2(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode | Exit Ethernet MA MEP mode. |
| PE2(config-ether-cfm-ma)#mep crosscheck mpid 20 | Configure cross-check to the remote MEP. |
| PE2(config-ether-cfm-ma)#exit-ether-ma-mode | Exit Ethernet MA mode. |
| PE2(config-ether-cfm)#exit | Exit Ethernet CFM mode and return to the configure mode. |
| PE2(config)#commit | Commit candidate configuration to be running configuration. |

Validation

The following output displays the validation results for PE1, PE2, and PE3 in an EVPN setup, which includes EVPN xconnect status, Ethernet CFM errors, remote maintenance points, local maintenance points, and successful ping tests.

PE1: Display xConnect Status

```
PE1#show evpn mpls xconnect
EVPN Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                               Remote      Connection-Details
=====
=====
VPN-ID      EVI-Name      MTU  VPN-ID      Source      Destination      PE-
IP          MTU   Type  NW-Status
=====
=====
52          ----          1500  2           xe33.2       00:00:00:aa:aa:bb:bb:00:00:00
7.7.7.7      1500  AC-NW  NW-SET
1500  ----  ----
```

PE1: Display Ethernet CFM Errors

```
PE1#show ethernet cfm errors domain evpn1

Domain Name      Level      MEPID      Defects
-----
evpn1            7          20         .....
```

PE1: Display Remote Maintenance Points

```
PE1#show ethernet cfm maintenance-points remote domain evpn1 ma-name evp1
MEPID      RMEPID      LEVEL      Rx CCM      RDI      PEER-MAC      TYPE
-----
20          10           7          Yes         False    00aa.bb00.0002 Configured
```

PE1: Display Local Maintenance Points

```
PE1#show ethernet cfm maintenance-points local mep domain evpn1 ma-name evp1
MPID Dir Lvl CC-Stat HW-Status CC-Intvl MAC-Address Def Port MD Name
-----
20 Up 7 Enable Installed 100 ms 3417.ebe4.af22 F xe33.2 evpn1
```

PE1: Ping Test

```
PE1#ping ethernet mac 00aa.bb00.0002 unicast source 20 domain evpn1 ma evp1
success rate is 100 (5/5)
```

PE2/PE3: Display xConnect Status

```
PE2#show evpn mpls xconnect
EVPN Xconnect Info
=====
AC-AC: Local-Cross-connect
AC-NW: Cross-connect to Network
AC-UP: Access-port is up
AC-DN: Access-port is down
NW-UP: Network is up
NW-DN: Network is down
NW-SET: Network and AC both are up

Local                               Remote      Connection-Details
=====
VPN-ID      EVI-Name      MTU  VPN-ID      Source      Destination      PE-
IP          MTU   Type   NW-Status
=====
2           1.1.1.1      ----  1500  52         po1.2          --- Single Homed Port ---
                1500  AC-NW  NW-SET
```

PE2/PE3: Display Ethernet CFM Errors

```
PE2#show ethernet cfm errors domain evpn1

Domain Name      Level      MEPID      Defects
-----
evpn1            7          10         .....
```

PE2/PE3: Display Local Maintenance Points

```
PE2#show ethernet cfm maintenance-points local mep domain evpn1 ma-name evp1
MPID Dir Lvl CC-Stat HW-Status CC-Intvl MAC-Address Def Port MD Name
-----
10 Up 7 Enable Installed 100 ms 00aa.bb00.0002 F po1.2 evpn1
```

PE2/PE3: Display Remote Maintenance Points

```
PE2#show ethernet cfm maintenance-points remote domain evpn1 ma-name evp1
```

| MEPID | RMEPID | LEVEL | Rx CCM | RDI | PEER-MAC | TYPE |
|-------|--------|-------|--------|-------|----------------|------------|
| 10 | 20 | 7 | Yes | False | 3417.ebe4.af22 | Configured |

PE2/PE3: Ping and Traceroute Test

```
PE2#ping ethernet mac 3417.ebe4.af22 unicast source 10 domain evpn1 ma evp1
success rate is 100 (5/5)
PE2#traceroute ethernet 3417.ebe4.af22 mepid 10 domain evpn1 ma evp1
MP Mac           Hops  Relay-action          Ingress/Egress  Ingress/Egress action
3417.ebe4.af22    1     RlyHit                Ingress         IngOK
```

Implementation Examples

Here is a practical scenario and use case for CFM (802.1ag) implementation for ELINE MultiHoming in the context of a telecommunications service provider who offers Ethernet-based Virtual Private Network (EVPN) services to various enterprises.

Use Case: Ensuring Service Reliability and Quality

Scenario: Consider a company with several branch offices that rely heavily on its ELINE connections to ensure smooth communication and data exchange between offices. The company subscribes to an EVPN service provided by a telecommunications service provider.

Use Case Details

- **MultiHoming Resilience:** The company's critical applications and services require high availability. MultiHoming ensures redundancy by connecting each branch office to the provider's network through multiple paths. This way, if one path fails due to network issues, the traffic can be rerouted through the alternative path without causing a disruption.
- **Continuous Monitoring:** CFM implementation allows the service provider to continuously monitor the connectivity and performance of the ELINE connections between the branch offices. By sending CCMs, the provider can quickly identify any interruptions in connectivity.
- **Swift Issue Detection and Resolution:** In case of a network disruption or fault, the service provider receives immediate alerts through CFM CCMs. This enables the provider's network operations team to pinpoint the issue's location and take prompt action to restore services, minimizing downtime for the company.
- **Troubleshooting Efficiency:** The CFM Ping and Trace functions assist in troubleshooting network issues. If the company's IT team reports a performance issue or communication problem, use CFM diagnostic capabilities to trace the path of packets and identify bottlenecks or faulty segments.

In this use case, the CFM implementation for ELINE MultiHoming provides a robust solution for ensuring reliable and high-quality connectivity for the company's distributed offices. It enables proactive monitoring and rapid issue resolution, which are critical for maintaining the company's communication and operational efficiency.

Troubleshooting

Follow the troubleshooting steps below to resolve connectivity issues related to CFM EVPN-ID and crosscheck local and remote MEP ID matching.

1. **Check Local EVPN-ID:** Verify the EVPN-ID configured on the local device (Example: PE1) and ensure that it matches the intended EVPN-ID for the target service or connection.

2. Verify Remote EVPN-ID: Check the EVPN-ID configured on the remote device (Example: PE2) and confirm that it matches the EVPN-ID expected by the local device.
3. Crosscheck MEP ID: Examine the MEP ID configured on the local device (PE1) and ensure it matches the expected R-MEP ID on the remote device (PE2).
4. Validate Remote MEP ID: Verify the MEP ID configured on the remote device (PE2) and ensure it matches the R-MEP ID expected by the local device (PE1).
5. Reconfigure If Needed: If there are discrepancies between the local and remote EVPN-IDs or MEP IDs, reconfigure the devices to match.
6. Test the Connection: After ensuring that EVPN-IDs and MEP IDs match on both devices, test the connection to confirm it is established correctly.

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

| Acronym | Description |
|--------------|----------------------------------|
| CFM | Continuity Fault Management |
| CCM | Continuity Check Messages |
| EVPN | Ethernet Virtual Private Network |
| ELINE | Ethernet-LINE |
| MPLS | Multi-Protocol Label Switching |
| UP | User-to-Provider |
| MEP | Maintenance End Point |
| MH | MultiHoming |
| R-MEP | Remote Maintenance End Point |
| MAC | Media Access Control |
| AC | Access Circuit |
| TLV | Type-Length-Value |
| RDI | Remote Defect Indicator |
| LB | Loopback |
| LBM | Loopback Message |
| LTR | Looptrace Reply |

Glossary

The following provides definitions for key terms used throughout this document.

| | |
|--|---|
| Continuity Fault Management (CFM) | A protocol (802.1ag) that facilitates the monitoring of network connectivity by using Continuity Check Messages (CCM) to detect faults. |
| Continuity Check Messages (CCM) | Periodic messages are used in CFM to check the continuity of a network connection. |

| | |
|--|---|
| Ethernet Virtual Private Network (EVPN) | A technology that enables the extension of Layer 2 Ethernet networks over a Layer 3 infrastructure. |
| Multi-Protocol Label Switching (MPLS) | In telecommunications networks, a routing technique is employed to guide data from one node to the next using concise path labels rather than relying on lengthy network addresses. |
| User-to-Provider (UP) | Refers to the connection between the user's and service provider's networks. |
| Maintenance End Point (MEP) | A point in a network where maintenance operations are initiated or terminated. |
| MultiHoming (MH) | A networking architecture that allows a device to be connected to multiple network paths or points of attachment. |
| Remote Maintenance End Point (R-MEP) | A maintenance end point located remotely from the originating point. |
| Media Access Control (MAC) | A unique identifier assigned to a network interface card. |
| Attachment Circuit (AC) | A circuit that connects a customer's equipment to a provider's network. |
| Type-Length-Value (TLV) | A format used to encapsulate data in a variety of protocols. |
| Remote Defect Indication (RDI) | A signal used to indicate a fault condition in a network link. |
| Loopback Message (LBM) | A message used in loopback testing to check connectivity. |

EVPN EPL Link-Loss Forwarding

Overview

Link Loss Forwarding (LLF) is a fault propagation feature for EVPN Ethernet Private Line (EPL) services that ensures reliable point-to-point connections on physical interfaces. It prevents traffic blackholing by detecting remote service failures and bringing the local physical link operationally down, enabling the local end to trigger a traffic failover mechanism. LLF is supported on all Broadcom chipsets and uses Ethernet Auto-Discovery (AD) per EVI (RT-1) route withdrawal for fault propagation, which is enabled by default.

The feature can be configured per service using the CLI command `llf enable` in interface mode. By facilitating proactive fault handling, LLF ensures minimal downtime, improves service reliability, and enhances traffic management for customers relying on point-to-point connections.

Feature Characteristics

Link Loss Forwarding (LLF) enhances the reliability of EVPN Ethernet Private Line (EPL) connections by detecting remote failures, bringing the local link down to initiate failover, and preventing traffic blackholing. It ensures prompt communication of service interruptions to minimize downtime and prevent data loss.

Benefits

LLF enhances networking environments by ensuring minimal service downtime, reliable traffic failover, and increased network stability.

- **Minimized Service Downtime:** Quick detection and propagation of faults reduce recovery time.
- **Traffic Failover Assurance:** Reliable failover mechanisms are enabled, ensuring minimal impact on customer services.
- **Improved Network Reliability:** The proactive nature of LLF enhances the robustness of EVPN EPL services.

Configuration of LLF for an EVPN Ethernet Private Line (EPL)


To configure LLF on an interface, follow the below steps to ensure that network faults are managed effectively, preventing traffic loss or blackholing by administratively adjusting the status of affected interfaces based on the fault condition.

Topology

The topology consists of two Customer Edge devices (CE1 and CE2) connected to Provider Edge devices (PE1 and PE2) through sub-interfaces. The Provider Edge devices are interconnected through MPLS. In this topology, the Link Loss Forwarding (LLF) configuration for an EVPN Ethernet Private Line (EPL) service ensures fault propagation to prevent traffic blackholing during link failures. When a link failure occurs, such as between CE-1 and PE-1 on interface xe1.1, PE-1 detects the issue and withdraws the Ethernet Auto-Discovery (AD) per EVI (RT-1) route. This withdrawal is propagated through the MPLS network to PE-2, which then takes action to bring down the corresponding physical interface (xe22.2). This mechanism ensures the fault is communicated end-to-end, allowing

CE-2 to detect the failure and trigger failover mechanisms. LLF is enabled through the LLF enable command at the interface level, ensuring automated fault detection and seamless service continuity.

Figure 63. Link-loss Forwarding Topology



Configuration XE22.2

Follow these steps to configure LLF on PE2.

1. Access the interface configuration mode, configure the interface xe22.2. Set the interface as a switch port.

```

PE2(config)#configure terminal
PE2(config)#int xe22.2 sw

```

2. Specify the encapsulation type as dot1q and set the VLAN ID to 10. Enable the interface to support EVPN services.

```

PE2(config-if)#encapsulation dot1q 10
PE2(config-if)#access-if-evpn

```

3. Enable LLF on the interface to ensure link failure detection and recovery. Save the changes and exit the configuration mode.

```

PE2(config-acc-if-evpn)#llf-enable
PE2(config)#commit
PE2(config)#exit

```

Validation

To validate the proper configuration of Link Loss Forwarding (LLF), use the `show running-config interface xe22.2` command to inspect the interface settings. Ensure that the interface is configured with switchport, encapsulation dot1q 10 for VLAN tagging, access-if-evpn to enable EVPN on the interface, llf-enable to activate LLF, and map vpn-id 10 to associate the service with the correct VPN ID. Verify that these attributes are present and correctly configured. If discrepancies are found, update the configuration and recheck to ensure the LLF setup is correct and operational.

```

#show running-config interface xe22.2
interface xe22.2 switchport
  encapsulation dot1q 10
  access-if-evpn
  llf-enable
  map vpn-id 10
!
Rvtep#sh evpn mpls llf-status
status : is_up\is_running

```

| ifp-name | sub-Ifname | llf-configured | RT-1-rcvd | llf-action |
|----------------|----------------|----------------|-----------|------------|
| xe16 (DOWN\--- | xe16.2 (UP\--- | YES | NO | YES |


```

Total number of entries are 1

Rvtep#sh int brief
xe16.2          SUBINTERFACE  --      --          down    PD      10g    --      No    No

Rvtep#sh int xe16
Interface xe16
  Flexport: Non Control Port (Active)
  Hardware is ETH  Current HW addr: e8c5.7a8f.c60e
  Physical:e8c5.7a8f.c60e  Logical:(not set)
  Forward Error Correction (FEC) configured is Auto (default)
  FEC status is N/A
  Port Mode is Router
  Protected Mode is Promiscuous
  Interface index: 10037
  GMPLS index: 39
  Metric 1 mtu 1500 duplex-full  link-speed 10g
  Debounce timer: disable
  ARP ageing timeout 1500
  <BROADCAST,MULTICAST>
  Reason: LLF DOWN
  VRF Binding: Not bound
  Label switching is disabled
  No Virtual Circuit configured
  Administrative Group(s): None
  Regular Extended-Admin-Group(s): None
  Anomaly Extended-Admin-Group(s): None
  Bandwidth 10g
  Maximum reservable bandwidth 10g
    Available b/w at priority 0 is 10g
    Available b/w at priority 1 is 10g
    Available b/w at priority 2 is 10g
    Available b/w at priority 3 is 10g
    Available b/w at priority 4 is 10g
    Available b/w at priority 5 is 10g
    Available b/w at priority 6 is 10g
    Available b/w at priority 7 is 10g
  DHCP client is disabled.
  Last Flapped: 2024 Oct 25 21:46:12 (00:00:44 ago)
  Load Interval: 30 seconds.
  Statistics last cleared: Never
  inet6 fe80::eac5:7aff:fe8f:c60e/64
  ND router advertisements are sent approximately every 442 seconds
  ND next router advertisement due in 0 seconds.
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses.
  RX
    unicast packets 0 multicast packets 0 broadcast packets 0
    input packets 0 bytes 0
    jumbo packets 0
    undersize 0 oversize 0 CRC 0 fragments 0 jabbers 0
    input error 0
    input with dribble 0 input discard 0
    Rx pause 0
  TX
    unicast packets 0 multicast packets 10 broadcast packets 0
    output packets 10 bytes 1020
    jumbo packets 0
    output errors 0 collision 0 deferred 0 late collision 0
    output discard 0
    Tx pause 0
Rvtep#

```

CLI Commands

The LLF introduces the following configuration commands in OcNOS.

llf-enable

Use this command to enable the link-loss forwarding on an access-interface.

Use the no parameter to disable the load interval globally.

Command Syntax

```
llf-enable
no llf-enable
```

Parameters

None

Default

Disabled

Command Mode

Access-if mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

The following example illustrates how to activate EVPN Ethernet Private Line (EPL).

```
#configure terminal
(config)#int xe22.2 sw
(config-if)#encapsulation dot1q 10
(config-if)#access-if-evpn
(config-acc-if-evpn)#llf-enable
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Table 87.

| Key Terms/Acronym | Description |
|-------------------|--|
| LLF | Link Loss Forwarding (LLF) is a feature in networking, particularly in Ethernet-based networks like Ethernet Virtual Private Network (EVPN), that is designed to prevent traffic blackholing during link failures. |
| EPL | Ethernet Private Line (EPL) is a high-performance, point-to-point Ethernet service used for dedicated connectivity between two locations. |
| MPLS | Multiprotocol Label Switching (MPLS) is a high-performance routing technique used in modern telecommunications networks to efficiently route data packets. |

TWAMP over EVPN Configuration

This chapter contains a complete sample TWAMP over EVPN configuration.

Two-way Active Measurement Protocol (TWAMP) is an open protocol for measuring network performance between any two devices. The TWAMP MPLS transport is implemented as part of supporting TWAMP on routers which acts as MPLS routers both in the roles of LERs as well as intermediate routers. Ocnos 6.0 also supports the end to end statistics calculation when multiple paths are available between sender and reflector with multihop support.

The user can use the link delay metrics such as average, minimum, and maximum delay, and delay variance to determine the network latency. Using link delay metrics will enable troubleshooting latency issues or apply Traffic Engineering (TE) solutions to meet Service Level Agreements (SLAs).

The TWAMP protocol is designed to do such measurements, and a basic implementation of this protocol has already been implemented in ocnos. This feature here is a TWAMP protocol in ocnos where the focus will be on accuracy and configurable advertisement of the measured data.

EVPN (based on MPLS) Supported scenarios:

In general, TWAMP over EVPN works on,

- CE-CE Overlay Only
- CE-PE Overlay Only
- PE-PE Both Under lay and over lay

Topology

The illustration displays a sample TWAMP over EVPN topology.

- CE1 and CE2 are customer edge routers
- PE1 and PE2 are IPv4 Provider Edge routers
- P1 is the router at the core of the IPv4 MPLS provider network




Figure 64. TWAMP over EVPN Topology

Configure TWAMP over EVPN for LDP

EVPN-Eline configurations

PE1

| | |
|--|--|
| #configure terminal | Enter Configure mode. |
| (config)# evpn mpls enable | Enable EVPN MPLS |
| (config)# evpn mpls vtep-ip-global 18.18.18.18 | Configuring VTEP global IP to loopback IP |
| (config)# mac vrf eline500 | Create a new mac VRF named eline100 |
| (config-vrf)# rd 20.20.20.20:100 | Assign the route distinguisher (RD) value as 20.20.20.20:100 |
| (config-vrf)# route-target both 111:111 | Configuring import and export value as 111:111 Support: route-target export route-target import |
| (config-vrf)#exit | Exit VRF mode |
| (config)#commit | Commit the configurations |
| 9(config)# evpn mpls id 500 xconnect tar-get-mpls-id 501 | Configure the EVPN-VPWS identifier with source identifier 2 and target identifier 501 |
| (con-fig-evpn-mpls)#host-reachability-protocol evpn-bgp eline500 | Mapping vrf "eline500" to EVPN-VPWS identifier |
| (config-evpn-mpls)#commit | Commit the transaction. |
| (config-evpn-mpls)#exit | Exit the EVPN MPLS mode and return to the configure mode. |
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 18.18.18.18 | Configure router id as loopback address |
| (config-router)#transport-address ipv4 18.18.18.18 | Configure ldp transport address as loopback address |
| (config-router)#exit | Exit from the router ldp mode |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 18.18.18.18/32 sec-ondary | Assign IP address to Loopback interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe8 | Enter Interface mode |
| (config-if)# ip address 10.1.1.18/24 | Assign IP address to interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe11.500 switchport | Creating L2 sub interface of physical interface xe11 |
| (config-if)# encapsulation dot1q 500 | Setting Encapsulation to dot1q with VLAN ID 2 |

| | |
|--|---|
| | Supported Encapsulation: dot1ad, dot1q, untagged, default |
| (config-if)# access-if-evpn | Entering Access mode for EVPN MPLS ID configuration |
| (config-acc-if-evpn)# map vpn-id 500 | Map vpn-id 500 to interface xe11.500 (VPWS) |
| (config-if)#exit | Exit interface mode |
| (config)# router ospf 1 | Enter Router OSPF mode |
| (config-router)# ospf router-id 18.18.18.18 | Configure OSPF router-id |
| (config-router)# network 10.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# network 18.18.18.18/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# commit | Commit the configurations |
| (config-router)# exit | Exit from router OSPF mode |
| (config)# router bgp 100 | Enter BGP router mode |
| (config-router)# bgp router-id 18.18.18.18 | Configure BGP router-id |
| (config-router)# neighbor 8.8.8.8 remote-as 100 | Configure PE2 as an iBGP4+ neighbor |
| (config-router)# neighbor 8.8.8.8 up-date-source lo | Update the source as loopback for iBGP peering with the remote PE2 router |
| (config-router)# address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)# neighbor 8.8.8.8 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exit form address family |
| (config-router)# commit | Commit the configurations |

P1

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 3.3.3.3 | Configure router id as loopback address |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure ldp transport address as loopback address |
| (config-router)#exit | Exit from the router ldp mode |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 3.3.3.3/32 secondary | Assign IP address to Loopback interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe14 | Enter Interface mode |
| (config-if)# ip address 10.1.1.3/24 | Assign IP address to interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |

| | |
|---|--|
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe15 | Enter Interface mode |
| (config-if)# ip address 11.1.1.3/24 | Assign IP address to interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)# router ospf 1 | Enter Router OSPF mode |
| (config-router)# ospf router-id 18.18.18.18 | Configure OSPF router-id |
| (config-router)# network 10.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# network 3.3.3.3/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# network 11.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# commit | Commit the configurations |
| (config-router)# exit | Exit from router OSPF mode |

PE2

| | |
|--|--|
| #configure terminal | Enter Configure mode. |
| (config)# evpn mpls enable | Enable EVPN MPLS |
| (config)# evpn mpls vtep-ip-global 8.8.8.8 | Configuring VTEP global IP to loopback IP |
| (config)# mac vrf eline500 | Create a new mac VRF named eline100 |
| (config-vrf)# rd 20.20.20.20:100 | Assign the route distinguisher (RD) value as 20.20.20.20:100 |
| (config-vrf)# route-target both 111:111 | Configuring import and export value as 111:111 Support: route-target export route-target import |
| (config-vrf)#exit | Exit VRF mode |
| (config)#commit | Commit the configurations |
| 9config)# evpn mpls id 501 xconnect tar-get-mpls-id 500 | Configure the EVPN-VPWS identifier with source identifier 501 and target identifier 500 |
| (con-fig-evpn-mpls)#host-reachability-protocol evpn-bgp eline500 | Mapping vrf "eline500" to EVPN-VPWS identifier |
| (config-evpn-mpls)#exit | Exit from evpn mpls mode |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 8.8.8.8/32 secondary | Assign IP address to Loopback interface |
| (config-if)#exit | Exit Interface mode |

| | |
|---|---|
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 8.8.8.8 | Configure router id as loopback address |
| (config-router)#transport-address ipv4 8.8.8.8 | Configure ldp transport address as loopback address |
| (config-router)#exit | Exit from the router ldp mode |
| (config)#interface xe24 | Enter Interface mode |
| (config-if)# ip address 11.1.1.8/24 | Assign IP address to interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe25 | Enter Interface mode |
| (config-if)# ip vrf forwarding vrf100 | Bind the interface connected to the CE1 router with VRF 100 |
| (config-if)# ip address 101.1.1.1/24 | Assign IP address to interface |
| (config-if)#exit | Exit interface mode |
| (config)# router ospf 1 | Enter Router OSPF mode |
| (config-router)# ospf router-id 8.8.8.8 | Configure OSPF router-id |
| (config-router)# network 11.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# network 8.8.8.8/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit from router OSPF mode |
| (config)# router bgp 100 | Enter BGP router mode |
| (config-router)# bgp router-id 8.8.8.8 | Configure BGP router-id |
| (config-router)# neighbor 18.18.18.18 re-mote-as 100 | Configure PE2 as an iBGP4+ neighbor |
| (config-router)# neighbor 18.18.18.18 up-date-source lo | Update the source as loopback for iBGP peering with the remote PE2 router |
| (config-router)# address-family l2vpn evpn | Entering into address family mode as EVPN |
| (config-router-af)# neighbor 8.8.8.8 activate | Enabling EVPN Address family for neighbor |
| (config-router-af)#exit | Exit form address family |
| (config-router)# commit | Commit the configurations |

CE1

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 37.37.37.37/32 secondary | Assign IP address to Loopback interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe24.500 | Enter Interface mode |

| | |
|--|--|
| (config-if)#encapsulation dot1q 500 | Setting Encapsulation to dot1q with VLAN ID 500 |
| (config-if)# ip address 172.16.10.1/24 | Configure IP address on the interface |
| (config-if)#exit | Exit Interface mode |
| (config)# router bgp 200 | Enter BGP router mode |
| (config-router)# bgp router-id 37.37.37.37 | Configure BGP router-id |
| (config-router)#neighbor 100.1.1.1 remote-as 100 | Configure PE1 as an eBGP4+ neighbor |
| (config-router)# address-family ipv4 unicast | Enter address-family IPv4 unicast mode |
| (config-router-af)# redistribute connected | Redistribute the connected route under address family IPv4 unicast |
| (config-router-af)# neighbor 100.1.1.1 activate | Activate the neighbor in the IPv4 address family |
| (config-router-af)#exit | Exit form address family |
| (config-router)# commit | Commit the configurations |

CE2

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 2.2.2.2/32 secondary | Assign IP address to Loopback interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe14.500 | Enter Interface mode |
| (config-if)#encapsulation dot1q 500 | Setting Encapsulation to dot1q with VLAN ID 500 |
| (config-if)# ip address 172.16.10.2/24 | Assign IP address to interface |
| (config-if)#exit | Exit Interface mode |
| (config)# router bgp 300 | Enter BGP router mode |
| (config-router)# bgp router-id 2.2.2.2 | Configure BGP router-id |
| (config-router)#neighbor 101.1.1.1 remote-as 100 | Configure PE2 as an eBGP4+ neighbor |
| (config-router)# address-family ipv4 unicast | Enter address-family IPv4 unicast mode |
| (config-router-af)# redistribute connected | Redistribute the connected route under address family |
| IPv4 unicast | |
| (config-router-af)# neighbor 101.1.1.1 activate | Activate the neighbor in the IPv4 address family |
| (config-router-af)#exit | Exit form address family |
| (config-router)# commit | Commit the configurations |

TWAMP Configuration on Sender (CE1)

TWAMP sender is configured to measure the delay on interface xe24.500 CE1

| | |
|---|--|
| #configure terminal | Enter Configure mode. |
| (config)# hardware-profile filter twamp-ipv4 enable | Enable hardware filter for ipv4 to configure TWAMP |

| | |
|---|---|
| | measure-ment configs |
| (config)#commit | Commit the configuration |
| (config)# twamp-light control | Enable TWAMP light controller on CE1 |
| (config-twamp-light-con)# control-admin-state enable | Enable TWAMP Controller admin state |
| (config)# interface xe24.500 | Enter Interface Loopback mode |
| (config-if)# delay-measurement dynamic twamp reflector-ip 172.16.10.2 | Configure delay measurement on interface to reflector CE2 |
| (config-if)#commit | Commit the configurations |
| (config-if)#end | Return to privilege mode |

TWAMP Configuration on Reflector (CE2)

Configure TWAMP Reflector as interface xe24.500 on CE2

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)# hardware-profile filter twamp-ipv4 enable | Enable hardware filter for ipv4 to configure TWAMP measure-ment configs |
| (config)#commit | Commit the configuration |
| (config)# twamp-light reflector | Enable TWAMP light Reflector on CE2 |
| (config-twamp-light-ref)# reflec-tor-admin-state enable | Enable the TWAMP reflector admin state |
| (config-twamp-light-ref)# reflector-name pe2 reflector-ip ipv4 172.16.10.2 | Configure TWAMP reflector IP as CE2 interface IP |
| (config-twamp-light-ref)#commit | Commit the configurations |
| (config-if)#end | Return to privilege mode |

Validation

1. Verify ping from CE1 to CE2

```
PE1#ping 172.16.10.2
Press CTRL+C to exit
PING 172.16.10.2 (172.16.10.2) 56(84) bytes of data.
 64 bytes from 172.16.10.2: icmp_seq=1 ttl=64 time=0.776 ms
 64 bytes from 172.16.10.2: icmp_seq=2 ttl=64 time=0.553 ms
```

2. Verify the TWAMP statistics on all the configured interfaces on CE1

In the below verification command, packets sent and received showing as equal. So all the TWAMP packets received reply for all the sent packets for the delay measurement. Showing all the Round Trip Delay and Reverse Delay timers.

```
PE1#sh twamp-statistics
=====
TWAMP Test-Session Statistics
=====
Test-Session Name      : __internal_interface_xe24.500
Start Time             : 2023 Mar 16 00:04:37
Elapsed time(milli sec) : 3001
```

```

Packets Sent           : 1
Packets Received       : 1
Packet Loss (%)        : 0.00
Round Trip Delay (usec)
  Minimum              : 383
  Maximum              : 383
  Average              : 383
Forward Delay (usec)
  Minimum              : (*)
  Maximum              : (*)
  Average              : (*)
Reverse Delay (usec)
  Minimum              : (*)
  Maximum              : (*)
  Average              : (*)
Round Trip Delay Variation (usec)
  Minimum              : 383
  Maximum              : 383
  Average              : 383
Forward Delay Variation (usec)
  Minimum              : (*)
  Maximum              : (*)
  Average              : (*)
Reverse Delay Variation (usec)
  Minimum              : (*)
  Maximum              : (*)
  Average              : (*)

```

(*) - Time is not in sync between Sender and Reflector

3. Verify the List of all interfaces that are currently participating in Delay measurement

```

PE1#sh twamp-statistics interfaces
Interface Last Advertisement Delay(us) Min(us) Max(us) Var(us) Loss(%)
xe24.500      -              0        0        0        0    Not Enabled

```

4. Verify the Detailed list of TWAMP delay measurement information on interface Loopback

```

PE1#sh twamp-statistics interfaces xe24.500
Interface name       : xe24.500
Sender IP           : 172.16.10.1
Reflector IP        : 172.16.10.2
Reflector port      : 862
DSCP value          : 0
HW Status           : HW rules installed
Last Advertised stats:
  Time: 2023-03-16 00:05:08
  Average delay      : 198
  Minimum delay      : 166
  Maximum delay      : 233
  Average delay variation: 5
  Minimum delay variation: 2
  Maximum delay variation: 10
  Packets sent       : 10
  Packets received    : 10
  Packets timeout     : 0
  Packet Loss: Not Enabled
Last Calculated stats:
  Time: 2023-03-16 00:05:08
  Average delay      : 198
  Minimum delay      : 166
  Maximum delay      : 233
  Average delay variation: 5
  Minimum delay variation: 2
  Maximum delay variation: 10
  Packets sent       : 10

```

```
Packets received      : 10
Packets timeout       : 0
Packet Loss : Not Enabled
```

LAYER 2 VIRTUAL PRIVATE NETWORK CONFIGURATION

| | |
|--|------|
| Virtual Private Wire Service Configuration | 2019 |
| Overview | 2019 |
| Configuration | 2019 |
| Sub-Interface with Multiple Match | 2028 |
| VPWS-CFM | 2030 |
| Virtual Private LAN Service Configuration | 2034 |
| VPLS Raw Mode | 2034 |
| VPLS Mesh Peer Configuration | 2034 |
| VPLS Tagged Mode | 2042 |
| VPLS Split Horizon Group | 2053 |
| VPLS - Per-Peer VPLS-type Signaling | 2068 |
| VPLS - Tunnel-ID | 2077 |
| VPLS - Tunnel-Name | 2087 |
| L2VPN FAT SUPPORT | 2100 |
| Topology | 2100 |
| Configuration for VPLS with FAT | 2100 |
| Configuration for BGP VPLS with FAT | 2111 |
| Configuration for VPWS with FAT | 2122 |
| L2VPN Sub Interface Configuration | 2132 |
| Topology | 2132 |
| VPLS Configuration | 2132 |
| VPWS Configuration | 2136 |
| Static VPLS Configuration | 2141 |
| Overview | 2141 |
| Configure Static VPLS | 2141 |
| Remove Configurations | 2148 |
| BGP-VPLS Configuration | 2149 |
| Topology | 2149 |
| BGP-VPLS Configuration | 2150 |
| BGP Four Bytes ASN with VPLS Configuration | 2162 |
| Static VPLS Service Mapping Configuration | 2174 |
| Overview | 2174 |
| Configuration | 2174 |
| Configuration of SubInterface | 2179 |
| LDP-VPLS Service Mapping Configuration | 2186 |

| | |
|---|------|
| Overview | 2186 |
| Configuration | 2186 |
| LDP VPLS Subinterface Matching | 2191 |
| BGP-VPLS Service Mapping Configuration | 2199 |
| Overview | 2199 |
| Configuration | 2199 |
| BGP VPLS SUBINTERFACE ATTACHMENT | 2204 |
| Hierarchical VPLS | 2213 |
| Overview | 2213 |
| Prerequisites | 2213 |
| Configuration for H-VPLS with Redundancy | 2214 |
| Configuration for H-VPLS without Redundancy | 2221 |
| Commands for H-VPLS Configuration | 2226 |
| MAC Withdrawal - VPLS/H-VPLS | 2229 |
| Overview | 2229 |
| Prerequisites | 2230 |
| Configuration | 2231 |
| Running Configuration | 2235 |
| Validation | 2237 |
| CLI Commands | 2240 |
| Glossary | 2243 |
| MAC Move Protection - VPLS/H-VPLS | 2244 |
| Overview | 2244 |
| Prerequisites | 2244 |
| Configuration | 2245 |
| Running Configuration | 2248 |
| Validation | 2249 |
| CLI Commands | 2252 |
| Glossary | 2255 |
| MAC Limit for VPLS and H-VPLS | 2256 |
| Overview | 2256 |
| Prerequisites | 2256 |
| Configuration | 2257 |
| Running Configuration | 2260 |
| Validation | 2261 |
| CLI Commands | 2263 |
| Glossary | 2265 |
| MLAG Active-Standby for VPLS | 2266 |
| Overview | 2266 |
| Prerequisites | 2267 |

| | |
|---|------|
| Configuration | 2267 |
| Running configurations | 2269 |
| Validation | 2270 |
| MLAG Active-Standby for VPLS commands | 2275 |
| Glossary | 2277 |
| BGP Auto-Discovery (AD) for LDP VPLS | 2279 |
| Overview | 2279 |
| Limitations | 2279 |
| Prerequisites | 2280 |
| Configuration | 2281 |
| Running Configuration | 2288 |
| Validation | 2294 |
| CLI Commands | 2304 |
| Ethernet Data Plane Loopback (EDPL) | 2310 |
| Overview | 2310 |
| Prerequisites | 2310 |
| Configuration | 2311 |
| Implementation Examples | 2315 |
| EDPL CLI Commands | 2317 |
| Troubleshooting | 2336 |
| Glossary | 2337 |

Virtual Private Wire Service Configuration


This chapter shows configurations for Virtual Private Wire Service (VPWS), where a point-to-point Layer 2 VPN service interconnects multiple Ethernet LANs across an MPLS backbone.

Overview

An MPLS Layer 2 Virtual Circuit (VC) is a point-to-point Layer 2 connection transported via MPLS on the service provider's network. The Layer 2 circuit is transported over a single Label Switched Path (LSP) tunnel between two Provider Edge (PE) routers.

The following diagram illustrates the configuration steps in this section. In this sample, the VC host devices, Host1 and Host2, are connected to the Provider Edge (PE) router PE-1; and Host3 and Host4 are connected to PE-2. The VC is established between PE-1 and PE-2. Interface xe2, on PE-1 and PE-2, is connected to the customer network; xe1, on PE-1 and PE-2, is connected to the MPLS cloud.

Figure 65. MPLS Layer 2 Virtual Circuit



The VC configuration process can be divided into the following steps:



Note: Loopback addresses being used should be advertised through OSPF, or should be statically routed.

1. Configure the IP address and OSPF for the PE-1, P (Provider), and PE-2 routers.
2. Configure MPLS and LDP on PE-1, P, and PE-2, and LDP targeted peer for the PE-1 and PE-2 routers. (If RSVP is used for configuring trunks, LDP must be configured on PE-1 and PE-2, and RSVP must be configured on PE-1, P, and PE-2.)
3. Configure the VC.
4. Bind the customer interface to the VC.

Configuration

Configure IP Address and OSPF on Routers

Configure the IP addresses and OSPF on the PE-1, P, and PE-2 routers.

PE-1

```
#configure terminal
```

Enter configure mode.

| | |
|---|---|
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface ce2/1 | Specify the interface (ce2/1) to be configured. |
| (config-if)#ip address 10.1.1.1/24 | Set the IP address of the interface to 10.1.1.1/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.1.1.0/24 area 0 (config-router)#network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

P

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 4.4.4.4/32 secondary | Set the IP address of the loopback interface to 4.4.4.4/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe11 | Specify the interface (xe11) to be configured. |
| (config-if)#ip address 10.1.1.2/24 | Set the IP address of the interface to 10.1.1.2/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe13 | Specify the interface (xe13) to be configured. |
| (config-if)#ip address 20.1.1.2/24 | Set the IP address of the interface to 20.1.1.2/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 10.1.1.0/24 area 0 (config-router)#network 20.1.1.0/24 area 0 (config-router)#network 4.4.4.4/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

PE-2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Specify the loopback interface (lo0) to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Specify the interface (xe15) to be configured. |
| (config-if)#ip address 20.1.1.2/24 | Set the IP address of the interface to 20.1.1.2/24. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#network 20.1.1.0/24 area 0 (config-router)#network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs, and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

Configure MPLS, LDP, and LDP Targeted Peer on Routers

Configure MPLS and LDP on PE-1, P, and PE-2, and LDP targeted peers on PE-1 and PE-2.



Note: If RSVP is used for configuring trunks, LDP must be configured on PE-1 and PE-2, and RSVP must be configured on PE-1, P, and PE-2,

PE-1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter the Router mode. |
| (config)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1 |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Specify the targeted LDP peer on PE-1. |
| (config-router-targeted-peer)# exit | Exit the Router targeted peer mode. |
| (config-router)#exit | Exit the Router mode. |
| (config)#interface ce2/1 | Specify the interface (xe1) to be configured. |
| (config-if)#label-switching | Enable label switching on interface ce2/1. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface ce2/1.. |
| (config-if)#commit | Commit the transaction. |

P

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter the Router mode. |
| (config)#router-id 4.4.4.4 | Set the router ID to IP address 4.4.4.4 |
| (config-router)#transport-address ipv4 4.4.4.4 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#exit | Exit the Router mode. |
| (config)#interface xe11 | Specify the interface (xe11) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe13. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe13 | Specify the interface (xe13) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe13. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe2. |
| (config-if)#commit | Commit the transaction. |

PE-2

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#router ldp | Enter the Router mode. |
| (config)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Specify the targeted LDP peer on PE-2. |
| (config-router-targeted-peer)# exit | Exit the Router targeted peer mode. |
| (config-router)#exit | Exit the Router mode. |
| (config)#interface xe15 | Specify the interface(xe15) to be configured. |
| (config-if)#label-switching | Enable label switching on interface xe15. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe15. |
| (config-if)#commit | Commit the transaction. |

Configure VC

Configure the VC. Each VC ID uniquely identifies the Layer-2 circuit among all the Layer-2 circuits.



Note: Both PE routers (endpoints) must be configured with the same VC-ID (100 in this example).

PE-1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#mpls l2-circuit vpws201 201 2.2.2.2 | Configure the VC for PE-2. In this example, vpws201 is the VC name, 201 is the VC ID, and 2.2.2.2 is the VC endpoint IP address. |
| (config)#mpls l2-circuit vpws200 200 2.2.2.2 | Configure the VC for PE-1 In this example, VPWS200 is the VC name. |
| (config)#commit | Save transaction into the database |

PE-2

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#mpls l2-circuit vpws202 202 1.1.1.1 | Configure the VC for PE-1. In this example, vpws202 is the VC name, 202 is the VC ID, and 1.1.1.1 is the VC endpoint IP address. |
| (config)#mpls l2-circuit vpws200 200 1.1.1.1 | Configure the VC for PE-2 In this example, VPWS200 is the VC name. |
| (config)#commit | Commit the transaction. |

Bind Customer Interface to VC

Bind the customer interface to the VC using one of the two procedures described below: Layer-2 untagged traffic or Layer-2 tagged traffic.



Note: Layer 2 VCs can only be bound to Layer 2 interfaces. The VC encapsulation method should be Ethernet (default), VLAN.

Layer 2 Untagged Traffic

Use Access mode for Layer 2 untagged traffic.

PE-1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface xe2.201 | Specify the interface (xe2.201) to be configured. |
| (config-if)#switchport | Switch to Layer-2 mode. |
| (config-if)#encapsulation default | Configure encapsulation to default. |
| (config-if)#access-if-vpws | Access VPWS under sub interface. |
| (config-acc-if-vpws)#mpls-l2-circuit vpws201 | Associating the VPWS Instance to the attachment circuit interface. |
| (config-if)#commit | Commit the transaction. |

PE-2

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface xe0.200 | Specify the interface (xe0.201) to be configured. |
| (config-if)#switchport | Switch to Layer-2 mode. |
| (config-if)#encapsulation default | Configure encapsulation to default. |
| (config-if)#access-if-vpws | Access VPWS under sub interface. |
| (config-acc-if-vpws)#mpls-l2-circuit vpws201 | Associating the VPWS Instance to the attachment circuit interface. |
| (config-if)#commit | Commit the transaction. |

Layer 2 Tagged Traffic

Use Trunk mode for Layer-2 tagged traffic. The following configuration allows only VLAN 2 and 3 traffic.

PE-1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#mpls l2-circuit vpws200 200 2.2.2.2 | Configure the VC for PE-2. In this example, vpws200 is the VC name, 200 is the VC ID, and 2.2.2.2 is the VC endpoint IP address. |
| (config)#interface xe2.200 | Specify the interface (xe2.200) to be configured. |
| (config-if)#switchport | Switch to Layer-2 mode. |
| (config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface. |
| (config-if)#access-if-vpws | Access VPWS under sub interface. |
| (config-acc-if-vpws)#mpls-l2-circuit vpws201 | Associating the VPWS Instance to the attachment circuit interface. |
| (config-if)#commit | Commit the transaction. |

PE-2

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#mpls l2-circuit t2 200 1.1.1.1 | Configure the VC for PE-2. In this example, t2 is the VC name, 200 is the VC ID, and 1.1.1.1 is the VC endpoint IP address. |
| (config-pseudowire)#exit | Exit pseudowire config mode. |
| (config)#interface xe0.200 | Specify the interface (xe0.200) to be configured. |
| (config-if)#switchport | Switch to Layer-2 mode. |
| (config-if)#access-if-vpws | Access VPWS under sub interface. |
| (config-acc-if-vpws)#mpls-l2-circuit vpws201 | Associating the VPWS Instance to the attachment circuit interface. |
| (config-if)#commit | Commit the transaction. |

Validation

Use the `show ldp mpls-l2-circuit` (Control Plane) command, and the `show mpls vc-table` (Forwarding Plane) command, to display complete information about the Layer 2 VC.

If the VC State is UP in the output from the `show ldp mpls-l2 circuit` command, and the Status is Active in the output of the `show mpls vc-table` command, a ping from CE1 to CE2 should be successful.

PE1

```
#show ldp mpls-l2-circuit
Transport Client VC VC Local Remote Destination Local
Remote
VC ID Binding State Type VC Label VC Label Address PW
Status PW Status
200 xe2.200 UP Ethernet VLAN
24961 24969 2.2.2.2 N/A N/A
201 xe2.100 UP Ethernet VLAN
24960 24971 2.2.2.2 N/A N/A

#show mpls vc-table
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VC-ID Access-Intf Network-Intf In label Out Label Tunnel-
Label Nexthop Status UpTime
200 xe2.200 ce2/1 24961 24969 24321 2.2.2.2 Active
01:57:25
201 xe2.100 ce2/1 24960 24971 24321 2.2.2.2 Active
01:43:59
```

PE2

```
PE2#show ldp mpls-l2-circuit
Transport Client VC VC Local Remote Destination Local
Remote
VC ID Binding State Type VC Label VC Label Address PW
Status PW Status
200 xe0.200 UP Ethernet VLAN
24969 24961 1.1.1.1 N/A N/A
201 xe0.201 UP Ethernet VLAN
24971 24960 1.1.1.1 N/A N/A

PE2#sh mpls vc-table
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VC-ID Access-Intf Network-Intf In label Out Label Tunnel-
Label Nexthop Status UpTime
200 xe0.200 xe15 24969 24961 24320 1.1.1.1 Active
01:59:08
201 xe0.201 xe15 24971 24960 24320 1.1.1.1 Active
01:45:42
```

These additional commands can also be used to display information about the Layer 2 virtual circuits.

```
show ldp mpls-l2-circuit detail
show ldp mpls-l2-circuit VC-ID
show ldp mpls-l2-circuit VC-ID detail
show mpls l2-circuit
```

PE1

```
PE1#show ldp mpls-l2-circuit detail
PW ID: 200, VC state is up
Access IF: xe2.200,up,AC state is up
Session IF: ce2/1, state is up
Destination: 2.2.2.2, Peer LDP Ident: 2.2.2.2
Local vctype: vlan, remote vctype: vlan
Local groupid: 0, remote groupid: 0
Local label: 24961, remote label: 24969
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: disabled Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : enabled
Current PW Status TLV : disabled
LDP-VC Signaled Time : 02:00:09

PW ID: 201, VC state is up
Access IF: xe2.201,up,AC state is up
Session IF: ce2/1, state is up
Destination: 2.2.2.2, Peer LDP Ident: 2.2.2.2
Local vctype: vlan, remote vctype: vlan
Local groupid: 0, remote groupid: 0
Local label: 24960, remote label: 24971
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: disabled Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : enabled
Current PW Status TLV : disabled
LDP-VC Signaled Time : 01:46:43
```

PE2

```
PE2#show ldp mpls-l2-circuit detail
PW ID: 200, VC state is up
Access IF: xe0.200,up,AC state is up
Session IF: xe15, state is up
Destination: 1.1.1.1, Peer LDP Ident: 1.1.1.1
Local vctype: vlan, remote vctype: vlan
Local groupid: 0, remote groupid: 0
Local label: 24969, remote label: 24961
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: disabled Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : enabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
LDP-VC Signaled Time : 02:01:19

PW ID: 201, VC state is up
Access IF: xe0.201,up,AC state is up
Session IF: xe15, state is up
Destination: 1.1.1.1, Peer LDP Ident: 1.1.1.1
Local vctype: vlan, remote vctype: vlan
Local groupid: 0, remote groupid: 0
Local label: 24971, remote label: 24960
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: disabled Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
```

```
Local PW Status Capability : enabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
LDP-VC Signaled Time : 01:47:53
```

Configure a Static Layer-2 VC

For a static MPLS Layer 2 VC configuration:

1. Configure the VC with the manual option
2. Configure the VC FIB entry
3. Bind the VC; all steps are in the configurations that follow.

PE-1

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| PE1(config)#mpls l2-circuit t3 300 2.2.2.2 | Configure the VC ID for PE1 |
| PE1(config-pseudowire)#manual-pseudowire | Configure pseudowire manual (no signaling) |
| PE1(config-pseudowire)#exit | Exit pseudowire config mode. |
| PE1(config)#interface xe2.300 | Enter sub interface mode |
| PE1(config)#switchport | config interface as switch port. |
| PE1(config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface. |
| PE1(config-if)#access-if-vpws | Access VPWS under sub interface. |
| PE1(config-acc-if-vpws)#mpls-l2-circuit vpws201 | Associating the VPWS Instance to the attachment circuit interface. |
| PE1(config-acc-if-vpw)#exit | Exit access mode. |
| PE1(config)#mpls l2-circuit-fib-entry 300 1000 2000 2.2.2.2 xe1 xe2.300 | Configure the VC ID with the manual option (no signaling used). |
| PE1(config)#commit | Commit the transaction. |

PE-2

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| PE2(config)#mpls l2-circuit t3 300 1.1.1.1 | Configure the VC ID for PE2 |
| PE2(config-pseudowire)#manual-pseudowire | Configure pseudowire manual (no signaling) |
| PE2(config)#interface xe0.300 | Enter sub interface mode |
| PE2(config)#switchport | Config interface as switch port. |
| PE2(config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface. |
| PE2(config-if)#access-if-vpws | Access VPWS under sub interface. |
| PE2(config-acc-if-vpws)#mpls-l2-circuit vpws201 | Associating the VPWS Instance to the attachment circuit interface. |
| PE2(config-acc-if-vpws)#exit | Exit access mode. |
| PE2(config)#mpls l2-circuit-fib-entry 300 2000 1000 1.1.1.1 xe1 xe0.300 | Configure the VC ID with the manual option (no |

| | |
|--------------------|-------------------------|
| | signaling used). |
| PE2(config)#end | Exit configure mode. |
| PE2(config)#commit | Commit the transaction. |

Validation

This example shows number of configured VCs and its status.

```
PE1#show mpls vc-table count
```

```
-----
Num PWS      : 3
Active PWS   : 3
OAM-only PWS : 0
Inactive PWS : 0
-----
```

```
PE1#show ldp mpls-l2-circuit count
```

```
-----
Num Signaled PWS: 3      [UP: 3]
-----
```

```
PE1#sh mpls vc-table
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VC-ID | Access-Intf | Network-Intf | In label | Out Label | Tunnel- | | |
|----------|-------------|--------------|----------|-----------|---------|---------|--------|
| Label | Nexthop | Status | UpTime | | | | |
| 300 | xe2.300 | N/A | 1000 | 2000 | 120 | 2.2.2.2 | Active |
| 00:00:05 | | | | | | | |

```
PE1#
```

```
PE2
```

```
PE2#sh mpls vc-table
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VC-ID | Access-Intf | Network-Intf | In label | Out Label | Tunnel- | | |
|----------|-------------|--------------|----------|-----------|---------|---------|--------|
| Label | Nexthop | Status | UpTime | | | | |
| 300 | xe0.300 | N/A | 2000 | 1000 | 150 | 1.1.1.1 | Active |
| 00:00:16 | | | | | | | |

Sub-Interface with Multiple Match

This is to validate the multiple match criteria support in a sub-interface. When multiple match statements are configured only rewrite push is supported, rewrite translate and pop are not supported.

PE-1

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#mpls l2-circuit t4 400 2.2.2.2 | Configure the VC for PE-1. In this example, t4 is the VC name, 400 is the VC ID, and 2.2.2.2 is the VC endpoint IP address. |
| (config-pseudowire)#exit | Exit pseudowire config mode. |

| | |
|---|--|
| (config)#interface xe2.400 | Enter sub interface mode |
| (config-if)#switchport | Switch to Layer-2 mode. |
| (config-if)# encapsulation dot1ad 3001 inner-dot1q 2001 | Configure encapsulation to dot1ad with outer VLAN ID 300 and inner VLAN ID 2001. |
| (config-if)#rewrite push dot1ad 401 | Configure encapsulation under a subinterface. |
| (config-if)# access-if-vpws | Access VPWS under sub interface. |
| (config-acc-if-vpws)# mpls-l2-circuit t4 primary | Configure the VC. In this example, t4 is the VC name |
| (config-acc-if-vpws)#commit | Commit the transaction. |

PE-2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#mpls l2-circuit vc1 10 1.1.1.1 | Configure the VC. |
| (config-pseudowire)#service-tpid dot1.ad | Configure Service-TPID as dot1.ad (0x88a8) |
| (config-pseudowire)#exit | Exit pseudowire config mode. |
| (config)#interface xe0.400 | Specify the interface (xe0.400) to be configured. |
| (config-if)#switchport | Switch to Layer-2 mode. |
| (config-if)# encapsulation dot1ad 3001 inner-dot1q 2001 | Configure encapsulation to dot1ad with outer VLAN ID 3001 and inner VLAN ID 2001. |
| (config-if)#rewrite push dot1ad 401 | Configure rewrite push. |
| (config-if)#commit | Commit the transaction. |
| (config-if)#end | End of interface and configurations mode. |

Validation

PE1

```
OcNOS#show ldp mpls-l2-circuit detail
PW ID: 400, VC state is up
Access IF: xe2.400,up,AC state is up
Session IF: ce2/1, state is up
Destination: 2.2.2.2, Peer LDP Ident: 2.2.2.2
Local vctype: vlan, remote vctype: vlan
Local groupid: 0, remote groupid: 0
Local label: 24962, remote label: 24970
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: disabled Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : disabled
Remote PW Status Capability : enabled
Current PW Status TLV : disabled
LDP-VC Signaled Time : 01:56:45

PE1#show mpls vc-table
OcNOS#sh mpls vc-table
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VC-ID | Access-Intf | Network-Intf | In label | Out Label | Tunnel- | | |
|----------|-------------|--------------|----------|-----------|---------|---------|--------|
| Label | Nexthop | Status | UpTime | | | | |
| 400 | xe2.400 | ce2/1 | 24962 | 24970 | 24321 | 2.2.2.2 | Active |
| 01:58:43 | | | | | | | |

PE2

```
PE2#show mpls vc-table
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VC-ID | Access-Intf | Network-Intf | In label | Out Label | Tunnel- | | |
|----------|-------------|--------------|----------|-----------|---------|---------|--------|
| Label | Nexthop | Status | UpTime | | | | |
| 400 | xe0.400 | xe15 | 24970 | 24962 | 24320 | 1.1.1.1 | Active |
| 01:59:52 | | | | | | | |

```
PE2# #sh ldp mpls-l2-circuit detail
PW ID: 201, VC state is up
Access IF: xe0.201,up,AC state is up
Session IF: xe15, state is up
Destination: 1.1.1.1, Peer LDP Ident: 1.1.1.1
Local vctype: vlan, remote vctype: vlan
Local groupid: 0, remote groupid: 0
Local label: 24971, remote label: 24960
Local MTU: 1500, Remote MTU: 1500
Local Control Word: disabled Remote Control Word: disabled Current use: disabled
Local Flow Label Direction: Disabled, Static: Disabled
Remote Flow Label Direction: Disabled, Static: Disabled
Local PW Status Capability : enabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
LDP-VC Signaled Time : 03:19:23
```

VPWS-CFM

The below examples contains of Ethernet Operations and Management (OAM) configurations using the Connectivity Fault Management (CFM) protocol.

Connectivity Fault Management detects, verifies, isolates and notifies connectivity failures on a Virtual Bridged LAN (B-VLAN) based on the protocol standard specified in IEEE 802.1ag 2007. It provides discovery and verification of paths through 802.1 bridges and LANs and is part of the Operation, Administration and Management (OAM) module. CFM is transparent to customer data being transported by a network and is capable of providing maximum fault management.

Prerequisite

Configure basic VPWS configuration as in above section.

1. Configure the IP address and OSPF for the PE-1, P (Provider), and PE-2 routers.
2. Configure MPLS and LDP on PE-1, P, and PE-2, and LDP targeted peer for the PE-1 and PE-2 routers. (If RSVP is used for configuring trunks, LDP must be configured on PE-1 and PE-2, and RSVP must be configured on PE-1, P, and PE-2.)

PE1

| | |
|---|---|
| PE1#configure terminal | Enter configure mode. |
| PE1(config)#ethernet cfm domain-type character-string domain-name MD-05 level 0 mip-creation none | Create cfm domain with type as character string and set mip creation criteria to default. |
| PE1(config-ether-cfm)#service ma-type string ma-name CFM-2 | Create ma type as string and configure the ma |
| PE1(config-ether-cfm-ma)# vpws vc1 | Configure vpws to associate to the MA |
| PE1(config-ether-cfm-ma)# ethernet cfm mep up mpid 115 active true vpws vc1 | Create mep up on vpws. |
| PE1(config-ether-cfm-ma-mep)#cc multicast state enable | Enable cc multicast. |
| PE1(config-ether-cfm-ma-mep)#exit-ether-ma-mep-mode | Exit CFM MEP configuration mode. |
| PE1(config-ether-cfm-ma)#mep crosscheck mpid 1115 | Configure crosscheck to remote MEP. |
| PE1(config-ether-cfm-ma)#cc interval 10ms | Enable cc interval for 10 millisecond. |
| PE1(config-ether-cfm-ma)#exit-ether-ma-mode | Exit CFM MA configuration mode. |
| PE1(config-ether-cfm)#exit | Exit ethernet CFM mode. |
| PE1(config)#commit | Commit the configuration |
| PE1(config)#exit | Exit the configure terminal mode |

PE2

| | |
|---|---|
| PE2#configure terminal | Enter configure mode. |
| PE2(config)#ethernet cfm domain-type character-string domain-name MD-05 level 0 mip-creation none | Create cfm domain with type as character string and set mip creation criteria to default. |
| PE2(config-ether-cfm)#service ma-type string ma-name CFM-2 | Create ma type as string and set mip creation criteria to default. |
| PE2(config-ether-cfm-ma)#vpws vc1 | Configure vpws to associate to the MA |
| PE2(config-ether-cfm-ma)#ethernet cfm mep down mpid 1115 active true vpws vc1 | Create mep up on vpws. |
| PE2(config-ether-cfm-ma-mep)#cc multicast state enable | Enable cc multicast. |
| PE2(config-ether-cfm-ma-mep)#exit-ether-ma- mep-mode | Exit CFM MEP configuration mode |
| PE2(config-ether-cfm-ma)#mep crosscheck mpid 115 | Configure crosscheck to remote MEP in VLAN 512. |
| PE2(config-ether-cfm-ma)#cc interval 10ms | Enable cc interval for 10 millisecond. |
| PE2(config-ether-cfm-ma)#exit-ether-ma-mode | Exit CFM MA configuration mode. |
| PE2(config-ether-cfm)#exit | Exit ethernet CFM mode. |
| PE2(config)#commit | Commit the configuration |
| PE2(config)#exit | Exit the configure terminal mode |

Validation

PE1

```

PE1#show ethernet cfm statistics
Continuity Check Messages
  CCM Sent      : 4378
  CCM Received   : 0

Loop Back Messages
  LBM Sent      : 5
  LBR Received(Valid) : 5
  LBR Received(Bad msdu) : 0
  LBR Received(Out-of-Seq): 0

Link Trace Messages
  LTM Sent      : 2
  LTR Sent      : 0
  LTR Received(Valid) : 2
  LTR Received(unexpected): 0

PE1#show ethernet cfm maintenance-points local mep domain MD-05 ma-name CFM-2
MPID Dir Lvl CC-Stat HW-Status CC-Intvl MAC-Address Def Port MD Name
-----
115 Up 0 Enable Installed 100 ms b86a.97db.2eca F xe4 MD-05

PE1#show ethernet cfm maintenance-points remote domain MD-05 ma-name CFM-2
MEPID RMEPID LEVEL Rx CCM RDI PEER-MAC TYPE
-----
115 1115 0 Yes False f88e.a192.4436 Configured

PE1#show ethernet cfm ma status domain MD-05 ma-name CFM-2
MA NAME STATUS
-----
CFM-2 Active

PE1#show ethernet cfm maintenance-points remote domain MD-05 ma-name CFM-2
MEPID RMEPID LEVEL Rx CCM RDI PEER-MAC TYPE
-----
115 1115 0 Yes False f88e.a192.4436 Configured

PE1#ping ethernet mac f88e.a192.4436 unicast source 115 domain MD-05 ma CFM-2
success rate is 100 (5/5)

PE1#traceroute ethernet f88e.a192.4436 mepid 115 domain MD-05 ma CFM-2
MP Mac Hops Relay-action Ingress/Egress Ingress/Egress action
f88e.a192.4436 1 RlyHit Ingress IngOK

```

PE2

```

PE2#show ethernet cfm statistics
Continuity Check Messages
  CCM Sent      : 8841
  CCM Received   : 0

Loop Back Messages
  LBM Sent      : 5
  LBR Received(Valid) : 5
  LBR Received(Bad msdu) : 0
  LBR Received(Out-of-Seq): 0

Link Trace Messages
  LTM Sent      : 1
  LTR Sent      : 2
  LTR Received(Valid) : 1
  LTR Received(unexpected): 0

```

```
PE2#show ethernet cfm maintenance-points local mep domain MD-05 ma-name CFM-2
```

| MPID | Dir | Lvl | CC-Stat | HW-Status | CC-Intvl | MAC-Address | Def Port | MD Name |
|------|-----|-----|---------|-----------|----------|----------------|----------|---------|
| 1115 | Up | 0 | Enable | Installed | 100 ms | f88e.a192.4436 | F xe19 | MD-05 |

```
PE2#show ethernet cfm maintenance-points remote domain MD-05 ma-name CFM-2
```

| MEPID | RMEPID | LEVEL | Rx CCM | RDI | PEER-MAC | TYPE |
|-------|--------|-------|--------|-------|----------------|------------|
| 1115 | 115 | 0 | Yes | False | b86a.97db.2eca | Configured |

```
PE2#show ethernet cfm ma status domain MD-05 ma-name CFM-2
```

| MA NAME | STATUS |
|---------|--------|
| CFM-2 | Active |

```
PE2#show ethernet cfm maintenance-points remote domain MD-05 ma-name CFM-2
```

| MEPID | RMEPID | LEVEL | Rx CCM | RDI | PEER-MAC | TYPE |
|-------|--------|-------|--------|-------|----------------|------------|
| 1115 | 115 | 0 | Yes | False | b86a.97db.2eca | Configured |

```
PE2#ping ethernet mac b86a.97db.2eca unicast source 1115 domain MD-05 ma CFM-2
success rate is 100 (5/5)
```

```
PE2#traceroute ethernet b86a.97db.2eca mepid 1115 domain MD-05 ma CFM-2
```

| MP Mac | Hops | Relay-action | Ingress/Egress | Ingress/Egress action |
|----------------|------|--------------|----------------|-----------------------|
| b86a.97db.2eca | 1 | RlyHit | Ingress | IngOK |

Virtual Private LAN Service Configuration


This chapter contains configurations for Virtual Private LAN Service (VPLS).

VPLS Raw Mode

Topology

The examples show the minimum configuration required for enabling a VPLS Mesh peer between PE1, PE2, and PE3 in raw mode.

Figure 66. VPLS Mesh Peers Raw Mode



VPLS Mesh Peer Configuration

PE1

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface ce2/1 | Specify the Interface (ce2/1) to be configured. |
| (config-if)#ip address 10.1.1.1/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID (100). |

| | |
|--|--|
| (config-router)#network 10.1.1.0/24 area 0 | Define the interface address on which the OSPF runs. |
| (config-router)#network 1.1.1.1/32 area 0 | Associate an area ID(0) with the interface address. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)# router-id 1.1.1.1 | Configure the router-id |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Configuring targeted LDP sessions to PE2 |
| (config-router-targeted-peer)#exit | Exit config-router-targeted-peer mode |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Enter targeted-peer-mode and PE3 |
| (config-router-targeted-peer)#exit | Exit config-router-targeted-peer mode |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface ce2/1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#mpls vpls VPLS1 100 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls-sig)#vpls-peer 2.2.2.2 (config-vpls-sig)#vpls-peer 3.3.3.3 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#interface xe2.100 switchport | Specify the attachment circuit interface. |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls VPLS1 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config-if)exit | Exit interface mode. |
| (config)#commit | Commit the transaction. |

P

| | |
|-----------------------|-----------------------|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |

| | |
|---|--|
| (config-if)#ip address 9.9.9.9/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe11 | Specify the interface to be configured. |
| (config-if)#ip address 10.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe13 | Enter interface mode. |
| (config-if)#ip address 20.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe12 | Enter interface mode. |
| (config-if)#ip address 30.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 10.1.1.0/24 area 0 (config-router)#network 20.1.1.0/24 area 0 (config-router)#network 30.1.1.0/24 area 0 (config-router)#network 4.4.4.4/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the Interface address. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 4.4.4.4 | Confirm the router-id |
| (config-router)#transport-address ipv4 4.4.4.4 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface configuration mode. |
| (config)#interface xe13 | Specify the interface to be configured. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe12 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#commit | Commit the transaction. |

PE2

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Specify the Interface (xe15) to be configured. |
| (config-if)#ip address 20.1.1.1/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bringing up the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 20.1.1.0/24 area 0 (config-router)#network 2.2.2.2/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 2.2.2.2 | Confirm the router-id |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configuring targeted LDP sessions to PE2. |
| (config-router-targeted-peer)#exit | Exit targeted-peer-mode |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configuring targeted LDP sessions to PE3 |
| (config-router-targeted-peer)#exit | Exit targeted-peer-mode |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#mpls vpls VPLS1 100 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls-sig)#vpls-peer 1.1.1.1 (config-vpls-sig)#vpls-peer 3.3.3.3 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)# interface xe0.100 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configuring the attachment circuit interface as |

| | |
|--|--|
| | Layer-2. |
| (config-if) #encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)# access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)# mpls-vpls VPLS1 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config-if)exit | Exit interface mode. |
| (config)#commit | Commit the transaction. |

PE3

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface cd6/1 | Specify the Interface (cd6/1) to be configured. |
| (config-if)#ip address 30.1.1.1/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively brining up the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 30.1.1.0/24 area 0 (config-router)#network 3.3.3.3/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 3.3.3.3 | Confirm the router-id |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configuring targeted LDP sessions to PE2 |
| (config-router-targeted-peer)#exit | Exit targeted-peer-mode |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Configuring targeted LDP sessions to PE3 |
| (config-router-targeted-peer)#exit | Exit targeted-peer-mode |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface cd6/1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |

| | |
|--|---|
| (config-if)#exit | Exit interface mode. |
| (config)#mpls vpls VPLS1 100 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls-sig)#vpls-peer 1.1.1.1 (config-vpls-sig)#vpls-peer 2.2.2.2 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)# int cd2/1.100 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configuring the attachment circuit interface as Layer-2. |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls VPLS1 | Associating the VPLS Instance to the attachment circuit interface |
| (config-if)#exit | Exit interface mode. |
| (config)#commit | Commit the transaction. |

Validation

PE1

Verify VPLS Session

```

PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
-----
2.2.2.2      ce2/1      Passive    OPERATIONAL  30      16:25:05
4.4.4.4      ce2/1      Passive    OPERATIONAL  30      16:25:02
3.3.3.3      ce2/1      Passive    OPERATIONAL  30      16:25:00

PE1#show mpls vpls detail
Virtual Private LAN Service Instance: VPLS1, ID: 100
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe2.100
Status: Up
Subinterface Match Criteria(s) :
```

```

dot1q 100

Mesh Peers:
  2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 13:49:20)
    FEC signaling element: FEC128
  3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 13:47:31)
    FEC signaling element: FEC128

PE1#sh mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID      Peer Addr      Tunnel-Label  In-Label  Network-Intf  Out-Label  Lkps/St  PW-INDEX
SIG-Protocol  Status          UpTime        Ext-Color
100          2.2.2.2         120           24964     ce2/1         24966      2/Up     6         L
DP           Active          13:50:08      -
100          3.3.3.3         103           24965     ce2/1         24960      2/Up     7         L
DP           Active          13:48:19      -

PE1#ping mpls vpls 100 peer 2.2.2.2/32 source 1.1.1.1
Sending 5 MPLS Echos to VPLS Id : 100, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

!
!
!
!
!

Success Rate is 100.00 percent (5/5)

```

P

Verify VPLS Mesh Peer

```

P#show ldp session
Codes: m - MD5 password is not set/unset.
      g - GR configuration not set/unset.
      t - TCP MSS not set/unset.
      Session has to be cleared manually

Code  Peer IP Address      IF Name  My Role  State      KeepAlive  UpTime
-----
1.1.1.1      xe11    Active   OPERATIONAL  30      16:25:35
2.2.2.2      xe13    Active   OPERATIONAL  30      16:25:41
3.3.3.3      xe12    Active   OPERATIONAL  30      16:41:51

```

PE3

```

PE2#show mpls vpls mesh

```

(m) - Service mapped over multipath transport
 (e) - Service mapped over LDP ECMP

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 100 | 1.1.1.1 | 24320 | 24961 | xe15 | 24962 | 2/Up | 2 L |
| DP | Active | 00:05:49 | - | | | | |
| 100 | 3.3.3.3 | 24321 | 24962 | xe15 | 24961 | 2/Up | 1 L |
| DP | Active | 00:05:25 | - | | | | |

PE2#show mpls vpls detail

Virtual Private LAN Service Instance: VPLS1, ID: 100
 SIG-Protocol: LDP
 Attachment-Circuit: UP
 Learning: Enabled
 Control-Word: Disabled
 Flow Label Status: Disabled, Direction: None, Static: No
 Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
 Description: none
 service-tpid: dot1q
 Operating mode: Raw
 MAC Withdrawal:

Configured interfaces:

Interface: xe0.100
 Status: Up
 Subinterface Match Criteria(s) :
 dot1q 100

Mesh Peers:

1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 13:52:01)
 FEC signaling element: FEC128
 3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 13:50:11)
 FEC signaling element: FEC128

PE2#show ldp session

Codes: m - MD5 password is not set/unset.
 g - GR configuration not set/unset.
 t - TCP MSS not set/unset.
 Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | xe15 | Active | OPERATIONAL | 30 | 00:14:03 |
| | 3.3.3.3 | xe15 | Passive | OPERATIONAL | 30 | 00:13:37 |
| | 4.4.4.4 | xe15 | Passive | OPERATIONAL | 30 | 00:14:04 |

PE3

PE3#show ldp session

Codes: m - MD5 password is not set/unset.
 g - GR configuration not set/unset.
 t - TCP MSS not set/unset.
 Session has to be cleared manually

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 4.4.4.4 | cd6/1 | Passive | OPERATIONAL | 30 | 00:09:01 |
| | 2.2.2.2 | cd6/1 | Active | OPERATIONAL | 30 | 00:09:25 |
| | 1.1.1.1 | cd6/1 | Active | OPERATIONAL | 30 | 00:09:15 |

PE3#show mpls vpls mesh

(m) - Service mapped over multipath transport
 (e) - Service mapped over LDP ECMP

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 100 | 1.1.1.1 | 24326 | 24961 | cd6/1 | 25601 | 2/Up | 6 |

| | | | | | | | | |
|-----|---------|----------|---|-------|-------|-------|------|---|
| LDP | Active | 00:09:33 | - | | | | | |
| 100 | 2.2.2.2 | 24327 | | 24960 | cd6/1 | 24960 | 2/Up | 5 |
| LDP | Active | 00:09:33 | - | | | | | |

```

PE3#show mpls vpls detail
Virtual Private LAN Service Instance: VPLS1, ID: 100
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

```

```

Configured interfaces:
Interface: cd2/1.100
Status: Up
Subinterface Match Criteria(s) :
dot1q 100

```

```

Mesh Peers:
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 13:52:43)
FEC signaling element: FEC128
2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 13:52:43)
FEC signaling element: FEC128

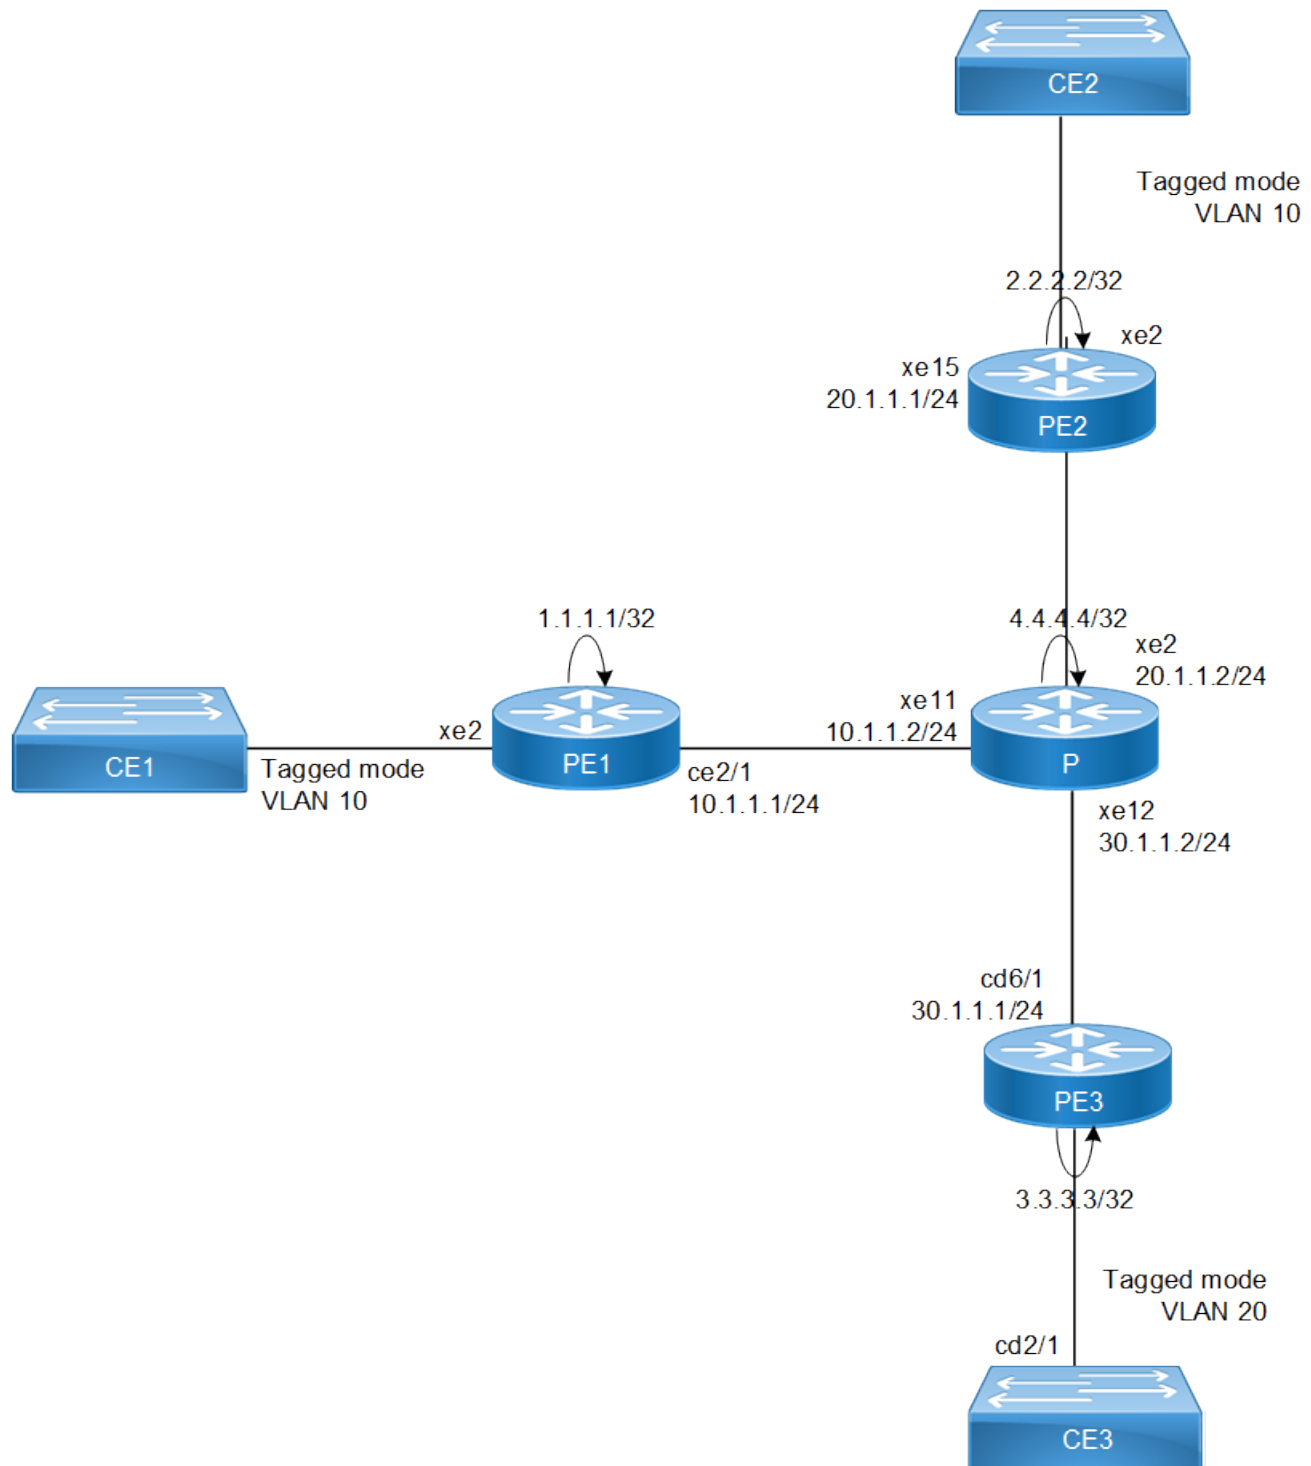
FEC signaling element: FEC128

```

VPLS Tagged Mode

Topology

The examples show the minimum configuration required for enabling a VPLS Mesh peer between PE1, PE2, and PE3 in Tagged Mode. In the below example PE1 and PE2 uses VLAN 10 for binding the VPLS instance to the attachment circuit and PE3 used VLAN 20 where it shows that VLAN swapping is supported.

Figure 67. VPLS Mesh Peers Tagged Mode

Configuration

PE1

```
#configure terminal
```

```
Enter configure mode.
```

| | |
|---|---|
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface ce2/1 | Specify the Interface (ce2/1) to be configured. |
| (config-if)#ip address 10.1.1.1/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively brining up the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 10.1.1.0/24 area 0 (config-router)#network 1.1.1.1/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Configuring targeted LDP sessions to PE2 |
| (config-router-targeted-peer)#exit | Exit targeted peer mode |
| (config-router-targeted-peer)#targeted-peer ipv4 3.3.3.3 | Configuring targeted LDP sessions to PE3 |
| (config-router-targeted-peer)#exit | Exit targeted peer mode |
| (config-router)#exit | Exit router configuration mode |
| (config)#interface ce2/1 | Enter interface mode |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#mpls vpls VPLS1 100 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls-sig)#vpls-type vlan | Configuring VPLS type as VLAN mode. |
| (config-vpls-sig)#vpls-peer 2.2.2.2 (config-vpls-sig)#vpls-peer 3.3.3.3 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#mpls vpls v4 28 | Enter VPLS config mode |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 2.2.2.2 | Configure VPLS Peers |
| config-vpls-sig)#vpls-peer 3.3.3.3 | Configure VPLS Peers |
| (config-vpls-sig)#exit-signaling | Exit Signaling LDP mode |

| | |
|--|---|
| (config-vpls)#exit | Exit VPLS mode |
| (config)#interface xe2.100 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configuring the attachment circuit interface as Layer-2. |
| (config-if)#encapsulation dot1q 100 | Configure the encapsulation |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if)# access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)# mpls-vpls VPLS1 | Associating the VPLS Instance to the attachment circuit interface |
| (config-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config)#interface xe2.28 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configuring the attachment circuit interface as Layer-2. |
| (config-if)#encapsulation dot1q 28 | Configure the encapsulation |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if)# access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)# mpls-vpls v4 | Associating the VPLS Instance to the attachment circuit interface |
| (config-if-vpls)#exit | Exit VPLS attachment-circuit mode |

P

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe11 | Specify the interface to be configured. |
| (config-if)#ip address 10.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe13 | Enter interface mode. |
| (config-if)#ip address 20.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe12 | Enter interface mode. |
| (config-if)#ip address 30.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 10.1.1.0/24 area 0 (config-router)#network 20.1.1.0/24 area 0 (config-router)#network 30.1.1.0/24 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the Interface |

| | |
|--|---|
| (config-router)#network 4.4.4.4/ 32 area 0 | address. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe11 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface configuration mode. |
| (config)#interface xe12 | Specify the interface to be configured. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe13 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#commit | Commit the transaction. |
| #copy running-config startup-config | Save the configuration. |

PE2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Specify the Interface (xe15) to be configured. |
| (config-if)#ip address 20.1.1.1/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively bring up the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 20.1.1.0/24 area 0 (config-router)#network 2.2.2.2/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configuring targeted LDP sessions to PE2 |
| (config-router-targeted-peer)#exit | Exit targeted peer mode |

| | |
|--|--|
| (config-router-targeted-peer)#targeted-peer ipv4 3.3.3.3 | |
| (config-router-targeted-peer)#exit | Exit targeted peer mode |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit targeted peer mode. |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#mpls vpls VPLS1 100 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls-sig)#vpls-type vlan | Configuring VPLS type as VLAN mode. |
| (config-vpls-sig)#vpls-peer 1.1.1.1 (config-vpls-sig)#vpls-peer 3.3.3.3 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#mpls vpls v4 28 | Enter VPLS config mode |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 1.1.1.1 | Configure VPLS Peers |
| config-vpls-sig)#vpls-peer 3.3.3.3 | Configure VPLS Peers |
| (config-vpls-sig)#exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#interface xe0.28 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configuring the attachment circuit interface as Layer-2. |
| (config-if)# encapsulation dot1q 28 | Configure encapsulation under a subinterface |
| (config-if)# access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)# mpls-vpls v4 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config)#interface xe0.100 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configuring the attachment circuit interface as Layer-2. |
| (config-if)# encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)# access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)# mpls-vpls VPLS1 | Associating the VPLS Instance to the attachment circuit interface. |

PE3

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface cd6/1 | Specify the Interface (cd6/1) to be configured. |
| (config-if)#ip address 30.1.1.1/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively brining up the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 30.1.1.0/24 area 0 (config-router)#network 3.3.3.3/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configuring targeted LDP sessions to PE2 |
| (config-router-targeted-peer)#exit | Exit targeted peer mode |
| (config-router-targeted-peer)#targeted-peer ipv4 2.2.2.2 | Configuring targeted LDP sessions to PE3 |
| (config-router-targeted-peer)#exit | Exit targeted peer mode |
| (config-router)#exit | Exit router configuration mode. |
| (config)#interface cd6/1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#mpls vpls VPLS1 100 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls-sig)#vpls-type vlan | Configuring VPLS type as VLAN mode. |
| (config-vpls-sig)#vpls-peer 1.1.1.1 (config-vpls-sig)#vpls-peer 2.2.2.2 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#mpls vpls v4 28 | Enter VPLS config mode |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 1.1.1.1 | Configure VPLS Peers |

| | |
|---------------------------------------|--|
| (config-vpls-sig)#vpls-peer 2.2.2.2 | Configure VPLS Peers |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#interface cd2/1.28 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configuring the attachment circuit interface as Layer-2. |
| (config-if)# encapsulation dot1q 28 | Configure encapsulation under a subinterface |
| (config-if)# access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)# mpls-vpls v4 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from subinterface |
| (config)#interface cd2/1.100 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configuring the attachment circuit interface as Layer-2. |
| (config-if)# encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)# access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)# mpls-vpls VPLS1 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit |

Validation

PE1

Verify VPLS Mesh Peer:

```
PE1#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 28 | 2.2.2.2 | 24322 | 24960 | ce2/1 | 24962 | 2/Up | 3 |
| DP | Active | 00:10:06 | - | | | | L |
| 28 | 3.3.3.3 | 24323 | 24962 | ce2/1 | 24960 | 2/Up | 4 |
| DP | Active | 00:13:16 | - | | | | L |
| 100 | 2.2.2.2 | 24322 | 24961 | ce2/1 | 24963 | 2/Up | 1 |
| DP | Active | 00:07:52 | - | | | | L |
| 100 | 3.3.3.3 | 24323 | 24963 | ce2/1 | 24961 | 2/Up | 2 |
| DP | Active | 00:13:16 | - | | | | L |

```
PE1#show mpls vpls detail
Virtual Private LAN Service Instance: v4, ID: 28
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Tagged
```

```

Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe2.28
Status: Up
Subinterface Match Criteria(s) :
dot1q 28

Mesh Peers:
2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:07:42)
FEC signaling element: FEC128
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:07:47)
FEC signaling element: FEC128

Virtual Private LAN Service Instance: VPLS1, ID: 100
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe2.100
Status: Up
Subinterface Match Criteria(s) :
dot1q 100

Mesh Peers:
2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:05:30)
FEC signaling element: FEC128
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:05:30)
FEC signaling element: FEC128

```

```
PE1#show ldp session
```

```

Codes: m - MD5 password is not set/unset.
g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|---------|-----------------|---------|-------------|-------|-----------|--------|
| 4.4.4.4 | ce2/1 | Passive | OPERATIONAL | 30 | 00:47:16 | |
| 2.2.2.2 | ce2/1 | Passive | OPERATIONAL | 30 | 00:10:07 | |
| 3.3.3.3 | ce2/1 | Passive | OPERATIONAL | 30 | 00:13:16 | |

P1

```
P#show ldp session
```

```

Codes: m - MD5 password is not set/unset.
g - GR configuration not set/unset.
t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|---------|-----------------|---------|-------------|-------|-----------|--------|
| 1.1.1.1 | xe11 | Active | OPERATIONAL | 30 | 00:53:29 | |
| 2.2.2.2 | xe13 | Active | OPERATIONAL | 30 | 00:16:24 | |
| 3.3.3.3 | xe12 | Active | OPERATIONAL | 30 | 00:54:39 | |

PE2

```

PE2#show mpls vpls detail
Virtual Private LAN Service Instance: v4, ID: 28
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe0.28
Status: Up
Subinterface Match Criteria(s) :
dot1q 28

Mesh Peers:
1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:07:42)
FEC signaling element: FEC128
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:07:47)
FEC signaling element: FEC128

Virtual Private LAN Service Instance: VPLS1, ID: 100
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe0.100
Status: Up
Subinterface Match Criteria(s) :
dot1q 100

Mesh Peers:
1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:05:30)
FEC signaling element: FEC128
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:05:30)
FEC signaling element: FEC128

PE2#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 28 | 1.1.1.1 | 24324 | 24962 | xe15 | 24960 | 2/Up | 1 |
| DP | Active | 00:07:57 | - | | | | L |
| 28 | 3.3.3.3 | 24325 | 24960 | xe15 | 24962 | 2/Up | 2 |
| DP | Active | 00:08:02 | - | | | | L |
| 100 | 1.1.1.1 | 24324 | 24963 | xe15 | 24961 | 2/Up | 3 |
| DP | Active | 00:05:45 | - | | | | L |

```

100      3.3.3.3      24325      24961      xe15      24963      2/Up      4      L
DP      Active      00:05:45      -

PE2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name      My Role      State      KeepAlive  UpTime
     4.4.4.4              xe15         Passive      OPERATIONAL 30      00:08:18
     1.1.1.1              xe15         Active       OPERATIONAL 30      00:08:17
     3.3.3.3              xe15         Passive      OPERATIONAL 30      00:08:18

```

PE3

```

PE3#show mpls vpls detail
Virtual Private LAN Service Instance: v4, ID: 28
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: cd4/2.28
Status: Up
Subinterface Match Criteria(s) :
dot1q 28

Mesh Peers:
  1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:38:57)
    FEC signaling element: FEC128
  2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:35:52)
    FEC signaling element: FEC128

Virtual Private LAN Service Instance: VPLS1, ID: 100
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: cd2/1.100
Status: Up
Subinterface Match Criteria(s) :
dot1q 100

Mesh Peers:
  1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:38:57)

```



```

FEC signaling element: FEC128
2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:33:35)
FEC signaling element: FEC128

PE3#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID      Peer Addr      Tunnel-Label  In-Label  Network-Intf  Out-Label  Lkps/St  PW-INDEX
SIG-Protocol  Status        UpTime      Ext-Color
28           1.1.1.1       24321       24960     cd6/1         24962      2/Up     2         L
DP           Active        00:39:24    -
28           2.2.2.2       24320       24962     cd6/1         24960      2/Up     1         L
DP           Active        00:36:19    -
100          1.1.1.1       24321       24961     cd6/1         24963      2/Up     4         L
DP           Active        00:39:24    -
100          2.2.2.2       24320       24963     cd6/1         24961      2/Up     3         L
DP           Active        00:34:02    -

PE3#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

Code  Peer IP Address      IF Name  My Role  State      KeepAlive  UpTime
4.4.4.4      cd6/1    Passive  OPERATIONAL  30    01:14:52
2.2.2.2      cd6/1    Active  OPERATIONAL  30    00:36:37
1.1.1.1      cd6/1    Active  OPERATIONAL  30    00:39:42

```

Validation for the Number of Configured VPLS Instances

This example below shows number of configured VPLS instances.

```

PE1#show mpls vpls count
-----
Total VPLS instances      : 2
Active VPLS instances     : 2
Inactive VPLS instances   : 2
-----

```

The example below shows the Count of VPLS from LDP standpoint

```

PE1#show ldp vpls count
-----
Total VPLS instances      : 2
Active VPLS instances     : 2
Inactive VPLS instances   : 0
-----

```

The example below shows the number of MAC addresses learned by VPLS.

```

PE1#show mpls vpls mac-address count
Total no of MAC addresses learnt :6

```

VPLS Split Horizon Group


In VPLS Hierarchy, multiple attachment circuits (AC) can be configured per few VPLS instances and AC-AC traffic blocking is not supported. Thus, when the PE router receives a broadcast, multicast, or unknown unicast packet on an AC. It sends the packet out on all other attachment circuits to all other CE devices participating in that VPLS instance.

To avoid this problem of a packet looping in the provider core, the PE devices enforce a "split-horizon" principle for the VPLS instances. That means the traffic will not flood if AC1 and AC2 are in same split horizon group.

Topology

The examples show the minimum configuration required for enabling a VPLS Mesh peer between PE-1, PE-2 with split horizon group. In the below example PE-1 and PE-2 uses split horizon groups where traffic can't be forwarded between the ACs if they are part of same access groups. But can send/receive from network ports.

Figure 68. Figure 8-15: VPLS Mesh Peers split horizon



Configuration With Service Template

PE1

| | |
|-----------------------|-----------------------|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |

| | |
|---|---|
| <code>(config-if)#ip address 4.4.4.4/32 secondary</code> | Configure IP address for the loopback interface. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe11</code> | Specify the Interface to be configured. |
| <code>(config-if)#ip address 100.1.1.1/24</code> | Configure IP address for the interface. |
| <code>(config-if)#no shutdown</code> | Administratively brining up the interface. |
| <code>(config-if)#label-switching</code> | Enable label switching on the interface |
| <code>(config-if)#mpls ldp-igp sync isis level-1-2</code> | Enable LDP IS-IS synchronization |
| <code>(config-if)#isis network point-to-point</code> | Configure the ISIS interface network type as point to point |
| <code>(config-if)#ip router isis ISIS-IGP</code> | Enable IS-IS routing on an interface |
| <code>(config-if)#enable-ldp ipv4</code> | Enable IPv4 LDP configuration on interface |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe13</code> | Specify the Interface to be configured |
| <code>(config-if)#ip address 100.3.1.1/24</code> | Configure IP address for the interface. |
| <code>(config-if)#no shutdown</code> | Administratively brining up the interface. |
| <code>(config-if)#label-switching</code> | Enable label switching on the interface |
| <code>(config-if)#mpls ldp-igp sync isis level-1-2</code> | Enable LDP IS-IS synchronization |
| <code>(config-if)#isis network point-to-point</code> | Configure the ISIS interface network type as point to point |
| <code>(config-if)#ip router isis ISIS-IGP</code> | Enable IS-IS routing on an interface |
| <code>(config-if)#enable-ldp ipv4</code> | Enable IPv4 LDP configuration on interface |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe12</code> | Specify the Interface to be configured |
| <code>(config-if)#ip address 100.2.1.1/24</code> | Configure IP address for the interface. |
| <code>(config-if)#no shutdown</code> | Administratively brining up the interface. |
| <code>(config-if)#label-switching</code> | Enable label switching on the interface |
| <code>(config-if)#mpls ldp-igp sync isis level-1-2</code> | Enable LDP IS-IS synchronization |
| <code>(config-if)#isis network point-to-point</code> | Configure the ISIS interface network type as point to point |
| <code>(config-if)#ip router isis ISIS-IGP</code> | Enable IS-IS routing on an interface |
| <code>(config-if)#enable-ldp ipv4</code> | Enable IPv4 LDP configuration on interface |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router isis ISIS-IGP</code> | Create an IS-IS routing instance |
| <code>(config)#is-type level-1</code> | Configure instance as level-2-only routing. |
| <code>(config)#metric-style wide</code> | Configure the new style of metric type as wide. |
| <code>(config)#mpls traffic-eng router-id 4.4.4.4</code> | Configure MPLS-TE unique router-id TLV. |

| | |
|---|--|
| (config)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1 |
| (config)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config)#fast-reroute ti-lfa level-1 proto ipv4 | Configure LFA-FRR to calculate the available backup path for all L1 IPv4 prefixes learned |
| (config)#fast-reroute ti-lfa level-2 proto ipv4 | Configure LFA-FRR to calculate the available backup path for all L2 IPv4 prefixes learned |
| (config)#bfd all-interfaces | Enable BFD for all neighbors. |
| (config)#net 49.0000.0000.0001.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config)#router ldp | Enter router mode for LDP. |
| (config)#router-id 4.4.4.4 | Enter router ID for LDP. |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit | Exit config-router-targeted-peer mode |
| (config-router)# transport-address ipv4 4.4.4.4 | Configuring targeted LDP sessions to PE-2 |
| (config-router)#exit | Exit router configuration mode. |
| (config)#router bgp 65010 | Define the routing process. The number 65010 specifies the AS number |
| (config-router)#bgp router-id 4.4.4.4 | Assign a BGP router ID. |
| (config-router)#neighbor 45.45.45.45 remote-as 65010 | Define the eBGP neighbor |
| (config-router)#neighbor 3.3.3.3 fall-over bfd multihop | Enabling BFD at BGP level for particular neighbor. |
| (config-router)#neighbor 3.3.3.3 update-source lo | Update the source for that particular neighbor as loopback interface |
| (config-router)#address-family ipv4 unicast | Enter into labeled-unicast address family |
| (config-router)#neighbor 3.3.3.3 activate | Activate the neighbor inside labeled-unicast address family |
| (config-router)#exit-address-family | Exit from address family IPv4 labeled unicast |
| (config-router)#address-family vpnv4 unicast | Enter into vpnv4 unicast address family |
| (config-router)#neighbor 3.3.3.3 activate | Activate the neighbor inside vpnv4 address family |
| (config-router)#exit-address-family | Exit from address family vpnv4 |
| (config-router)#address-family ipv6 unicast | Enter into labeled-unicast address family |
| (config-router)#redistribute connected | Redistribute the connected routes. |
| (config-router)#exit-address-family | Exit from address family IPv6 labeled unicast |
| (config)#mpls vpls VPLS-PE1-PE2-600 600 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |

| | |
|---|--|
| (config-vpls)#vpls-type vlan | Configure VPLS type as VLAN encapsulation |
| (config-vpls-sig)#vpls-peer 3.3.3.3 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#mpls vpls VPLS-PE1-PE2-601 601 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls)#vpls-type vlan | Configure VPLS type as VLAN encapsulation |
| (config-vpls-sig)#vpls-peer 3.3.3.3 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#mpls vpls VPLS-PE1-PE2-602 602 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls)#vpls-type vlan | Configure VPLS type as VLAN encapsulation |
| (config-vpls-sig)#vpls-peer 3.3.3.3 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#interface xe8.600 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configuring the attachment circuit interface as Layer-2. |
| (config-if)# encapsulation dot1q 600 | Configure the encapsulation |
| (config-if)# access-if-vpls | Associating the VPLS Instance to the attachment circuit interface. |
| (config-if)#mpls-vpls VPLS-PE1-PE2-600 | Attach the VPLS instance |
| (config-if)#split-horizon group access1 | Configuring split horizon group. |
| (config-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config)#interface xe8.601 | Create the subinterface |
| (config-if)#switchport | Associating the VPLS Instance to the attachment circuit interface. |
| (config-if)# encapsulation dot1q 601 | Configure the encapsulation |
| (config-if)# access-if-vpls | Configuring split horizon group |
| (config-if)#mpls-vpls VPLS-PE1-PE2-601 | Exit VPLS attachment-circuit mode |
| (config-if)#split-horizon group access1 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-if-vpls)#exit | Configuring split horizon group |
| (config)#interface xe8.602 | Create the subinterface |

| | |
|---|--|
| (config-if)#switchport | Exit VPLS attachment-circuit mode |
| (config-if)# encapsulation dot1q 601 | Configure the encapsulation |
| (config-if)# access-if-vpls | Specify the attachment circuit interface. |
| (config-if)#mpls-vpls VPLS-PE1-PE2-602 | Configuring the attachment circuit interface as Layer-2. |
| (config-if)#split-horizon group access1 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-if-vpls)#exit | Configuring split horizon group |

P1

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface ce2/1 | Specify the interface to be configured. |
| (config-if)#ip address 100.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enabling label switching capability on router |
| (config-if)#mpls ldp-igp sync isis level-1-2 | Enable LDP IS-IS synchronization |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface |
| (config-if)#enable-ldp ipv4 | Enable IPv4 LDP configuration on interface |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe12 | Specify the Interface to be configured |
| (config-if)#ip address 30.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#no shutdown | Administratively brining up the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#isis network point-to-point | Enable LDP IS-IS synchronization |
| (config-if)#ip router isis ISIS-IGP | Configure the ISIS interface network type as point to point |
| (config-if)#enable-ldp ipv4 | Enable IS-IS routing on an interface |
| (config-if)#exit | Enable IPv4 LDP configuration on interface |
| (config)#interface ce2/2 | Exit interface mode. |
| (config-if)#ip address 100.5.1.1/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#mpls ldp-igp sync isis level-1-2 | Enable LDP IS-IS synchronization |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |

| | |
|---|--|
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enable IPv4 LDP configuration on interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router isis ISIS-IGP | Create an IS-IS routing instance |
| (config)#is-type level-1 | Configure instance as level-2-only routing. |
| (config)#metric-style wide | Configure the new style of metric type as wide. |
| (config)#mpls traffic-eng router-id 1.1.1.1 | Configure MPLS-TE unique router-id TLV. |
| (config)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1 |
| (config)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config)#fast-reroute ti-lfa level-1 proto ipv4 | Configure LFA-FRR to calculate the available backup path for all L1 IPv4 prefixes learned |
| (config)#fast-reroute ti-lfa level-2 proto ipv4 | Configure LFA-FRR to calculate the available backup path for all L2 IPv4 prefixes learned |
| (config)#fast-reroute ti-lfa level-2 proto ipv4 | Configure LFA-FRR to calculate the available backup path for all L2 IPv4 prefixes learned |
| (config)#bfd all-interfaces | Enable BFD for all neighbors. |
| (config)#net 49.0000.0000.0002.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#exit | Exit router mode. |
| (config-router)# router-id 1.1.1.1 | Configure the router-id |
| (config-router)# transport-address ipv4 1.1.1.1 | Configure the transport address |
| (config-router)#exit | Exit router mode. |
| (config)#commit | Commit the transaction. |
| #copy running-config startup-config | Save the configuration. |

PE2

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#ip router isis ISIS-IGP | Attaching the ISIS |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Specify the Interface to be configured. |
| (config-if)#ip address 100.3.1.2/24 | Configure IP address for the interface. |

| | |
|--|--|
| <code>(config-if)#no shutdown</code> | Administratively brining up the interface. |
| <code>(config-if)#label-switching</code> | Enable label switching on the interface |
| <code>(config-if)#mpls ldp-igp sync isis level-1-2</code> | Enable LDP IS-IS synchronization |
| <code>(config-if)#isis network point-to-point</code> | Configure the ISIS interface network type as point to point |
| <code>(config-if)#ip router isis ISIS-IGP</code> | Enable IS-IS routing on an interface |
| <code>(config-if)#enable-ldp ipv4</code> | Enable IPv4 LDP configuration on interface |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe0</code> | Specify the Interface to be configured. |
| <code>(config-if)#ip address 100.5.1.2/24</code> | Configure IP address for the interface. |
| <code>(config-if)#no shutdown</code> | Administratively brining up the interface. |
| <code>(config-if)#label-switching</code> | Enable label switching on the interface |
| <code>(config-if)#mpls ldp-igp sync isis level-1-2</code> | Enable LDP IS-IS synchronization |
| <code>(config-if)#isis network point-to-point</code> | Configure the ISIS interface network type as point to point |
| <code>(config-if)#ip router isis ISIS-IGP</code> | Enable IS-IS routing on an interface |
| <code>(config-if)#enable-ldp ipv4</code> | Enable IPv4 LDP configuration on interface |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#interface xe14</code> | Specify the Interface to be configured. |
| <code>(config-if)#ip address 100.6.1.1/24</code> | Configure IP address for the interface. |
| <code>(config-if)#no shutdown</code> | Administratively brining up the interface. |
| <code>(config-if)#label-switching</code> | Enable label switching on the interface |
| <code>(config-if)#mpls ldp-igp sync isis level-1-2</code> | Enable LDP IS-IS synchronization |
| <code>(config-if)#isis network point-to-point</code> | Configure the ISIS interface network type as point to point |
| <code>(config-if)#ip router isis ISIS-IGP</code> | Enable IS-IS routing on an interface |
| <code>(config-if)#enable-ldp ipv4</code> | Enable IPv4 LDP configuration on interface |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router isis ISIS-IGP</code> | Create an IS-IS routing instance |
| <code>(config)#is-type level-1</code> | Configure instance as level-2-only routing. |
| <code>(config)#metric-style wide</code> | Configure the new style of metric type as wide. |
| <code>(config)#mpls traffic-eng router-id 3.3.3.3</code> | Configure MPLS-TE unique router-id TLV. |
| <code>(config)#mpls traffic-eng level-1</code> | Enable MPLS-TE in is-type Level-1 |
| <code>(config)#capability cspf</code> | Enable CSPF feature for ISIS instance. |
| <code>(config)#dynamic-hostname</code> | Configure the hostname to be advertised for an ISIS instance |
| <code>(config)#fast-reroute ti-lfa level-1 proto ipv4</code> | Configure LFA-FRR to calculate the available |

| | |
|--|--|
| | backup path for all L1 IPv4 prefixes learned |
| (config)#fast-reroute ti-lfa level-2 proto ipv4 | Configure LFA-FRR to calculate the available backup path for all L2 IPv4 prefixes learned |
| (config)#bfd all-interfaces | Enable BFD for all neighbors. |
| (config)#net 49.0000.0000.0003.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configuring targeted LDP sessions to PE-2 |
| (config-router)#pw-status-tlv | Configure the pw-status-tlv |
| (config-router)#targeted-peer ipv4 4.4.4.4 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit | Exit config-router-targeted-peer mode |
| (config-router-targeted-peer)#transport-address ipv4 3.3.3.3 | Configure the transport address |
| (config-router-targeted-peer)#exit | Exit config-router-targeted-peer mode |
| (config-router)#exit | Exit router configuration mode. |
| (config)#router bgp 65010 | Define the routing process. The number 65010 specifies the AS number |
| (config-router)#bgp router-id 3.3.3.3 | Assign a BGP router ID. |
| (config-router)#neighbor 4.4.4.4 remote-as 65010 | Define the eBGP neighbor |
| (config-router)#neighbor 4.4.4.4 fall-over bfd multihop | Enabling BFD at BGP level for particular neighbor. |
| (config-router)#neighbor 4.4.4.4 update-source lo | Update the source for that particular neighbor as loopback interface |
| (config-router)#address-family ipv4 unicast | Enter into labeled-unicast address family |
| (config-router)#neighbor 4.4.4.4 activate | Activate the neighbor inside labeled-unicast address family |
| (config-router)#exit-address-family | Exit from address family IPv4 labeled unicast |
| (config-router)#address-family vpnv4 unicast | Enter into vpnv4 unicast address family |
| (config-router)#neighbor 48.48.48.48 activate | Activate the neighbor inside vpnv4 address family |
| (config-router)#exit-address-family | Exit from address family vpnv4 |
| (config-router)#address-family ipv6 unicast | Enter into labeled-unicast address family |
| (config-router)#redistribute connected | Redistribute the connected routes. |
| (config-router)#exit-address-family | Exit from address family IPv6 labeled unicast |
| (config)#mpls vpls VPLS-PE1-PE2-600 600 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls)#vpls-type vlan | Configure VPLS type as VLAN encapsulation |
| (config-vpls-sig)#vpls-peer 4.4.4.4 | Configuring VPLS mesh peers. |

| | |
|--|--|
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#mpls vpls VPLS-PE1-PE2-601 601 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls)#vpls-type vlan | Configure VPLS type as VLAN encapsulation |
| (config-vpls-sig)#vpls-peer 4.4.4.4 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#mpls vpls VPLS-PE1-PE2-602 602 | Configuring VPLS instance with name and VPLS ID. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance. |
| (config-vpls)#vpls-type vlan | Configure VPLS type as VLAN encapsulation |
| (config-vpls-sig)#vpls-peer 4.4.4.4 | Configuring VPLS mesh peers. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode. |
| (config-vpls)#exit | Exit from VPLS Mode. |
| (config)#interface xe12.600 | Specify the attachment circuit interface. |
| (config-if)#switchport | Configure the switchport |
| (config-if)# encapsulation dot1q 600 | Configure the encapsulation |
| (config-if)# access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls VPLS-PE1-PE2-600 | Attach the VPLS instance |
| (config-acc-if-vpls)#split-horizon group access1 | Configuring split horizon group. |
| (config-acc-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe12.601 | Create the subinterface |
| (config-if)#switchport | Configure the switchport |
| (config-if)# encapsulation dot1q 601 | Configure the encapsulation |
| (config-if)# access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls VPLS-PE1-PE2-601 | Attach the VPLS instance |
| (config-acc-if-vpls)#split-horizon group access1 | Configuring split horizon group |
| (config-acc-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config-if)#exit | Exit the interface mode. |
| (config)#interface xe12.602 | Create the subinterface |
| (config-if)#switchport | Configure the switchport |
| (config-if)# encapsulation dot1q 601 | Configure the encapsulation |
| (config-if)# access-if-vpls | Access VPLS under sub interface |

| | |
|--|-----------------------------------|
| (config-acc-if-vpls)#mpls-vpls VPLS-PE1-PE2-602 | Attach the VPLS instance |
| (config-acc-if-vpls)#split-horizon group access1 | Configuring split horizon group |
| (config-acc-if-vpls)#exit | Exit VPLS attachment-circuit mode |
| (config-if)#exit | Exit the interface mode. |

P2

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#ip router isis ISIS-IGP | Attach the ISIS |
| (config-if)#exit | Exit interface mode. |
| (config)#interface ce6/1 | Specify the interface to be configured. |
| (config-if)#ip address 100.2.1.2/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#mpls ldp-igp sync isis level-1-2 | Enable LDP IS-IS synchronization |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enable IPv4 LDP configuration on interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface cd4/2 | Enter interface mode. |
| (config-if)#ip address 100.4.1.2/24 | Specify the interface to be configured. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#mpls ldp-igp sync isis level-1-2 | Enable LDP IS-IS synchronization |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enable IPv4 LDP configuration on interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface cd27/2 | Enter interface mode. |
| (config-if)#ip address 100.6.1.2/24 | Specify the interface to be configured. |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#mpls ldp-igp sync isis level-1-2 | Enable LDP IS-IS synchronization |
| (config-if)#isis network point-to-point | Configure the ISIS interface network type as point to point |
| (config-if)#ip router isis ISIS-IGP | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enable IPv4 LDP configuration on interface. |

| | |
|---|---|
| (config-if)#exit | Exit interface mode. |
| (config)#router isis ISIS-IGP | Create an IS-IS routing instance |
| (config)#is-type level-1 | Configure instance as level-2-only routing. |
| (config)#metric-style wide | Configure the new style of metric type as wide. |
| (config)#mpls traffic-eng router-id 2.2.2.2 | Configure MPLS-TE unique router-id TLV. |
| (config)#mpls traffic-eng level-1 | Enable MPLS-TE in is-type Level-1 |
| (config)#capability cspf | Enable CSPF feature for ISIS instance. |
| (config)#dynamic-hostname | Configure the hostname to be advertised for an ISIS instance |
| (config)#fast-reroute ti-lfa level-1 proto ipv4 | Configure LFA-FRR to calculate the available backup path for all L1 IPv4 prefixes learned |
| (config)#fast-reroute ti-lfa level-2 proto ipv4 | Configure LFA-FRR to calculate the available backup path for all L2 IPv4 prefixes learned |
| (config)#bfd all-interfaces | Enable BFD for all neighbors. |
| (config)#net 49.0000.0000.0004.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID. |
| (config-router)#exit | Exit router mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 2.2.2.2 | Configure router ID. |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#exit | Exit router mode. |
| (config)#commit | Commit the transaction. |
| #copy running-config startup-config | Save the configuration. |

Validation

PE1

```

PE1#sh ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

Code  Peer IP Address      IF Name   My Role   State      KeepAlive  UpTime
-----
1.1.1.1      xe11      Active    OPERATIONAL  30      01:09:57
2.2.2.2      xe12      Active    OPERATIONAL  30      00:50:08
3.3.3.3      xe13      Active    OPERATIONAL  30      01:00:15

PE1#show clns neighbor

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag ISIS-IGP:  VRF : default
System Id      Interface  SNPA              State Holdtime  Type Protocol

```

```

0000.0000.0002 xe11      0090.fb7d.ad0f      Up      20      L1      IS-IS
0000.0000.0004 xe12      0e00.0000.0019      Up      20      L1      IS-IS
0000.0000.0003 xe13      b86a.97c6.33d5      Up      26      L1      IS-IS

```

```
PE1#sh mpls vpls mesh
```

```
(m) - Service mapped over multipath transport
```

```
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 600 | 3.3.3.3 | 3 | 24961 | xe13 | 24960 | 2/Up | 2 L |
| DP | Active | 00:58:37 | - | | | | |
| 601 | 3.3.3.3 | 3 | 24962 | xe13 | 24961 | 2/Up | 1 L |
| DP | Active | 00:30:14 | - | | | | |
| 602 | 3.3.3.3 | 3 | 24963 | xe13 | 24962 | 2/Up | 3 L |
| DP | Active | 00:30:14 | - | | | | |

```
PE1#show mpls vpls detail
```

```
Virtual Private LAN Service Instance: VPLS-PE1-PE2-600, ID: 600
```

```
SIG-Protocol: LDP
```

```
Attachment-Circuit: UP
```

```
Learning: Enabled
```

```
Control-Word: Disabled
```

```
Flow Label Status: Disabled, Direction: None, Static: No
```

```
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
```

```
Description: none
```

```
service-tpid: dot1q
```

```
Operating mode: Tagged
```

```
Svlan Id: 0
```

```
Svlan Tpid: 8100
```

```
MAC Withdrawal:
```

```
Configured interfaces:
```

```
Interface: xe8.600
```

```
Status: Up
```

```
Subinterface Match Criteria(s) :
```

```
dot1q 600
```

```
Split-horizon group : access1
```

```
Mesh Peers:
```

```
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:58:51)
```

```
FEC signaling element: FEC128
```

```
Virtual Private LAN Service Instance: VPLS-PE1-PE2-601, ID: 601
```

```
SIG-Protocol: LDP
```

```
Attachment-Circuit: UP
```

```
Learning: Enabled
```

```
Control-Word: Disabled
```

```
Flow Label Status: Disabled, Direction: None, Static: No
```

```
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
```

```
Description: none
```

```
service-tpid: dot1q
```

```
Operating mode: Tagged
```

```
Svlan Id: 0
```

```
Svlan Tpid: 8100
```

```
MAC Withdrawal:
```

```
Configured interfaces:
```

```
Interface: xe8.601
```

```
Status: Up
```

```
Subinterface Match Criteria(s) :
```

```
dot1q 601
```

```
Split-horizon group : access1
```

```
Mesh Peers:
```

```
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:30:28)
```

```
FEC signaling element: FEC128
```

```

Virtual Private LAN Service Instance: VPLS-PE1-PE2-602, ID: 602
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe8.602
Status: Up
Subinterface Match Criteria(s) :
dot1q 602
Split-horizon group : access1

Mesh Peers:
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:30:28)
FEC signaling element: FEC128

PE2:

PE2#show ldp session
Codes: m - MD5 password is not set/unset.
      g - GR configuration not set/unset.
      t - TCP MSS not set/unset.
      Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
-----
4.4.4.4      xe15      Passive    OPERATIONAL  30      01:02:03
1.1.1.1      xe0       Active     OPERATIONAL  30      01:04:32
2.2.2.2      xe14      Active     OPERATIONAL  30      00:25:12

PE2#show clns neighbor

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag ISIS-IGP: VRF : default
System Id      Interface  SNPA              State  Holdtime  Type  Protocol
-----
0000.0000.0002 xe0        0090.fb7d.ad2a    Up     20        L1   IS-IS
0000.0000.0004 xe14       0e00.0000.0028    Up     25        L1   IS-IS
0000.0000.0001 xe15       0090.fb7d.ad21    Up     26        L1   IS-IS

PE2#sh mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID  Peer Addr      Tunnel-Label  In-Label  Network-Intf  Out-Label  Lkps/St  PW-INDEX
-----
SIG-Protocol  Status      UpTime      Ext-Color
600          4.4.4.4      3           24960     xe15          24961      2/Up     7         L
DP           Active      00:59:47    -
601          4.4.4.4      3           24961     xe15          24962      2/Up     8         L
DP           Active      00:31:23    -
602          4.4.4.4      3           24962     xe15          24963      2/Up     9         L
DP           Active      00:31:23    -

PE2#Show mpls vpls detail

```

```
Virtual Private LAN Service Instance: VPLS-PE1-PE2-600, ID: 600
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe12.600
Status: Up
Subinterface Match Criteria(s) :
dot1q 600
Split-horizon group : access1

Mesh Peers:
4.4.4.4 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 01:00:00)
FEC signaling element: FEC128

Virtual Private LAN Service Instance: VPLS-PE1-PE2-601, ID: 601
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe12.601
Status: Up
Subinterface Match Criteria(s) :
dot1q 601
Split-horizon group : access1

Mesh Peers:
4.4.4.4 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:31:36)
FEC signaling element: FEC128

Virtual Private LAN Service Instance: VPLS-PE1-PE2-602, ID: 602
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe12.602
```

```
Status: Up
Subinterface Match Criteria(s) :
dot1q 602
Split-horizon group : access1
```

```
Mesh Peers:
4.4.4.4 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:31:36)
FEC signaling element: FEC128
```

VPLS - Per-Peer VPLS-type Signaling

Overview

Virtual Private LAN Service (VPLS) LDP signaling is a mechanism used to establish and manage pseudowires (PWs) in a VPLS network. VPLS allows geographically dispersed sites to share a common Ethernet broadcast domain, enabling any-to-any communication between the sites as if they were all connected to the same local area network (LAN).

VPLS LDP signaling occurs when each PE has discovered the endpoints of the VPLS instance. PWs are established over MPLS tunnels between VPN sites to transparently transmit Layer 2 packets. With the ability to configure the VPLS type per peer, and with the VPLS type configured being forwarded to LDP only when there is a change in the VPLS type at the VPLS instance or per peer level, signaling behavior is more efficient. Previously, the VPLS type at the VPLS instance was significant across all PEs involved in the VPLS instance.


If there is a change in the VPLS type at the VPLS instance or per peer level, NSM will update LDP, triggering LDP signaling. If there is no change, there is no need for re-signaling, and the VPLS session remains in the same state. The VPLS type configured on PEs at both ends of a pseudowire must match; otherwise, negotiation is required, and the VPLS session will be inactive until resolved.

Feature Characteristics

Provide the ability to configure the VPLS type on a per-peer basis, in addition to the VPLS type set across all PEs in the VPLS instance. The VPLS type configured on PEs at both ends of a pseudowire must match; otherwise, negotiation should occur. The per-peer VPLS type should take priority when both the VPLS instance type and the per-peer type are configured.

Topology

The examples show the minimum configuration required for enabling a VPLS Mesh peer between PE1, PE2, and PE3 in per peer VPLS-type Signaling.

Figure 69. VPLS Per Peer VPLS Type Signaling

Configuration

The VPLS BGPSignaling feature enables you to use BGP as the control plane protocol for both auto discovery and signaling in accordance with RFC 4761

PE1 Configuration

1. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```
PE1(config)#configure terminal
PE1(config)#interface lo
PE1(config-if)#ip address 1.1.1.1/32 secondary
PE1(config-if)#exit
PE1(config)#interface ce2/1
PE1(config-if)#ip address 10.1.1.1/24
PE1(config-if)#no shutdown
PE1(config-if)#exit
PE1(config)#router ospf 100
PE1(config-router)#network 10.1.1.0/24 area 0
PE1(config-router)#network 1.1.1.1/32 area 0
PE1(config-router)#exit
PE1(config)#router ldp
PE1(config-router)# router-id 1.1.1.1
PE1(config-router)#transport-address ipv4 1.1.1.1
PE1(config-router)#targeted-peer ipv4 2.2.2.2
PE1(config-router-targeted-peer)#exit
PE1(config-router)#targeted-peer ipv4 3.3.3.3
PE1(config-router-targeted-peer)#exit
PE1(config-router)#exit
PE1(config)#interface ce2/1
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)#exit
PE1(config)#mpls vpls VPLS1 100
```

```
PE1(config-vpls)#signaling ldp
PE1(config-vpls-sig)#vpls-peer 2.2.2.2 vpls-type vlan
PE1(config-vpls-sig)#vpls-peer 3.3.3.3 vpls-type ethernet
PE1(config-vpls-sig)#exit-signaling
PE1(config-vpls)#exit
PE1(config)#int xe2.101 switchport
PE1(config-if)#encapsulation dot1q 1
PE1(config-if)#access-if-vpls
PE1(config-acc-if-vpls)#mpls-vpls VPLS101
PE1(config-acc-if-vpls)#commit
```

2. To exit the configuration, execute the following command.

```
PE1(config-if-vpls)#exit
PE1(config-if)#exit
```

3. To save the configuration, execute the following command.

```
PE1(config)#commit
PE1(config)#copy running-config startup-config
```

P Configuration

1. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```
P(config)#configure terminal
P(config)#interface lo
P(config-if)#ip address 4.4.4.4/32 secondary
P(config-if)#exit
P(config)#interface xe1
P(config-if)#ip address 10.1.1.2/24
P(config-if)#exit
P(config)#interface xe2
P(config-if)#ip address 20.1.1.2/24
P(config-if)#exit
P(config)#interface xe3
P(config-if)#ip address 30.1.1.2/24
P(config-if)#exit
P(config)#router ospf 100
P(config-router)#network 10.1.1.0/24 area 0
P(config-router)#network 20.1.1.0/24 area 0
P(config-router)#network 30.1.1.0/24 area 0
P(config-router)#network 4.4.4.4/32 area 0
P(config-router)#exit
P(config)#router ldp
P(config-router)#router-id 4.4.4.4
P(config-router)#transport-address ipv4 4.4.4.4
P(config-router-targeted-peer)#exit
P(config-router)#exit
P(config)#interface xe11
P(config-if)#label-switching
P(config-if)#enable-ldp ipv4
P(config-if)#exit
P(config)#interface xe12
P(config-if)#label-switching
P(config-if)#enable-ldp ipv4
P(config-if)#exit
P(config)#interface xe13
P(config-if)#label-switching
P(config-if)#enable-ldp ipv4
```

2. To exit the configuration, execute the following command.

```
P(config-if)#exit
```

3. To save the configuration, execute the following command.

```
P(config)#commit
P(config)#copy running-config startup-config
```

PE2 Configuration

1. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```
PE2(config)#configure terminal
PE2(config)#interface lo
PE2(config-if)#ip address 2.2.2.2/32 secondary
PE2(config-if)#exit
PE2(config)#interface xe15
PE2(config-if)#ip address 20.1.1.1/24
PE2(config-if)#no shutdown
PE2(config-if)#exit
PE2(config)#router ospf 100
PE2(config-router)#network 20.1.1.0/24 area 0
PE2(config-router)#network 2.2.2.2/32 area 0
PE2(config-router)#exit
PE2(config)#router ldp
PE2(config-router)#router-id 2.2.2.2
PE2(config-router)#transport-address ipv4 2.2.2.2
PE2(config-router)#targeted-peer ipv4 1.1.1.1
PE2(config-router-targeted-peer)#exit
PE2(config-router)#targeted-peer ipv4 3.3.3.3
PE2(config-router-targeted-peer)#exit
PE2(config-router)#exit
PE2(config)#interface xe15
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)#exit
PE2(config)#mpls vpls VPLS1 100
PE2(config-vpls)#signaling ldp
PE2(config-vpls-sig)#vpls-peer 2.2.2.2 vpls-type vlan
PE2(config-vpls-sig)#vpls-peer 3.3.3.3 vpls-type ethernet
PE2(config-vpls-sig)#exit-signaling
PE2(config-vpls)#exit
PE2(config)#int xe0.101 switchport
PE2(config-if)#encapsulation dot1q 1
PE2(config-if)#access-if-vpls
PE2(config-acc-if-vpls)#mpls-vpls VPLS101
PE2(config-acc-if-vpls)#commit
```

2. To exit the configuration, execute the following command.

```
PE2(config-if-vpls)#exit
PE2(config-if)#exit
```

3. To save the configuration, execute the following command.

```
PE2(config)#commit
PE2(config)#copy running-config startup-config
```

PE3 Configuration

1. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```
PE3(config)#configure terminal
PE3(config)#interface lo
PE3(config-if)#ip address 3.3.3.3/32 secondary
PE3(config-if)#exit
PE3(config)#interface cd6/1
PE3(config-if)#ip address 30.30.1.1/24
```

```

PE3(config-if)#no shutdown
PE3(config-if)#exit
PE3(config)#router ospf 100
PE3(config-router)#network 30.1.1.0/24 area 0
PE3(config-router)#network 3.3.3.3/32 area 0
PE3(config-router)#exit
PE3(config)#router ldp
PE3(config-router)#router-id 3.3.3.3
PE3(config-router)#transport-address ipv4 3.3.3.3
PE3(config-router)#targeted-peer ipv4 1.1.1.1
PE3(config-router-targeted-peer)#exit
PE3(config-router)#targeted-peer ipv4 2.2.2.2
PE3(config-router-targeted-peer)#exit
PE3(config-router)#exit
PE3(config)#interface ce6/1
PE3(config-if)#label-switching
PE3(config-if)#enable-ldp ipv4
PE3(config-if)#exit
PE3(config)#mpls vpls VPLS1 100
PE3(config-vpls)#signaling ldp
PE3(config-vpls-sig)#vpls-peer 2.2.2.2 vpls-type vlan
PE3(config-vpls-sig)#vpls-peer 3.3.3.3 vpls-type ethernet
PE3(config-vpls-sig)#exit-signaling
PE3(config-vpls)#exit
PE3(config)#interface cd2/1.101 switchport
PE3(config-if)# encapsulation dot1q 1
PE3(config-if)# access-if-vpls
PE3(config-acc-if-vpls)# mpls-vpls VPLS101
PE3(config-acc-if-vpls)#commit

```

2. To exit the configuration, execute the following command.

```

PE3(config-if-vpls)#exit
PE3(config-if)#exit

```

3. To save the configuration, execute the following command.

```

PE3(config)#commit
PE3(config)#copy running-config startup-config

```

Validation

PE1

```

PE1#show mpls vpls detail
Virtual Private LAN Service Instance: VPLS101, ID: 101
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe2.101
Status: Up
Subinterface Match Criteria(s) :
dot1q 1

Mesh Peers:
  2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d04h58m)
    FEC signaling element: FEC128
  3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d04h03m)

```

FEC signaling element: FEC128

```
PE1#show mpls vpls VPLS101
Virtual Private LAN Service Instance: VPLS1, ID: 100
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw

Configured interfaces:
Interface: xe0
Status: Up
Split-horizon group : access1
Service-template : st1
Match criteria : Accept all

Mesh Peers:
  2.2.2.2 (Peer VPLS Type: Ethernet VLAN) (Up) (UpTime: 00:55:45)
  3.3.3.3 (Peer VPLS Type: Ethernet) (Up) (UpTime: 00:55:01)

PE1#show ldp vpls 101
VPLS Identifier      : 101
Peer IP              : 2.2.2.2
VC State             : UP
VC Type              : vlan
VC Label Sent        : 24960
VC Label Received    : 24972
Local MTU             : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
Local FEC Type       : 128
Remote FEC Type      : 128
LDP-VPLS Signaled Time : 1d05h00m

VPLS Identifier      : 101
Peer IP              : 3.3.3.3
VC State             : UP
VC Type              : ethernet
VC Label Sent        : 24961
VC Label Received    : 24961
Local MTU             : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
```

```

Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
Local FEC Type       : 128
Remote FEC Type      : 128
LDP-VPLS Signaled Time : 1d04h07m

```

PE1#show mpls vpls mesh

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | SIG- |
|----------|-----------|--------------|-----------|--------------|-----------|---------|----------|------|
| Protocol | Status | UpTime | Ext-Color | | | | | |
| 101 | 2.2.2.2 | 120 | 24960 | ce2/1 | 24972 | 2/Up | 11 | L |
| DP | Active | 1d05h01m | - | | | | | |
| 101 | 3.3.3.3 | 103 | 24961 | ce2/1 | 24961 | 2/Up | 12 | L |
| DP | Active | 1d04h05m | - | | | | | |

PE1#show ldp session

```

Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | ce2/1 | Passive | OPERATIONAL | 30 | 1d05h20m |
| | 3.3.3.3 | ce2/1 | Passive | OPERATIONAL | 30 | 1d04h09m |
| | 4.4.4.4 | ce2/1 | Passive | OPERATIONAL | 30 | 1d05h20m |

PE1#show ip ospf neighbor

Total number of full neighbors: 1

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 4.4.4.4 | 1 | Full/Backup | 00:00:35 | 10.1.1.2 | ce2/1 | 0 |

PE2

PE2#show mpls vpls mesh

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 101 | 1.1.1.1 | 201 | 24972 | xe15 | 24960 | 2/Up | 15 |
| DP | Active | 1d05h03m | - | | | | |
| 101 | 3.3.3.3 | 301 | 24973 | xe15 | 24960 | 2/Up | 16 |
| DP | Active | 1d04h08m | - | | | | |

PE2#show mpls vpls VPLS101

```

Virtual Private LAN Service Instance: VPLS101, ID: 101
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
  Interface: xe0.101

```

```

Status: Up
Subinterface Match Criteria(s) :
dot1q 1

Mesh Peers:
  1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d05h04m)
    FEC signaling element: FEC128
  3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d04h09m)
    FEC signaling element: FEC128

PE2#show ldp vpls 101
VPLS Identifier      : 101
Peer IP              : 1.1.1.1
VC State             : UP
VC Type              : vlan
VC Label Sent        : 24972
VC Label Received    : 24960
Local MTU             : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
LDP-VPLS Signaled Time : 1d05h13m

VPLS Identifier      : 101
Peer IP              : 3.3.3.3
VC State             : UP
VC Type              : vlan
VC Label Sent        : 24973
VC Label Received    : 24960
Local MTU             : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
LDP-VPLS Signaled Time : 1d04h20m

PE2#show mpls vpls mesh

PE2#show ldp session
Peer IP Address      IF Name  My Role  State      KeepAlive UpTime

```

| | | | | | |
|---------|------|---------|-------------|----|----------|
| 1.1.1.1 | xe15 | Active | OPERATIONAL | 30 | 00:41:59 |
| 3.3.3.3 | xe15 | Passive | OPERATIONAL | 30 | 00:30:15 |
| 4.4.4.4 | xe15 | Passive | OPERATIONAL | 30 | 00:42:14 |

PE3

```
PE3#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | cd6/1 | Active | OPERATIONAL | 30 | 1d04h22m |
| | 2.2.2.2 | cd6/1 | Active | OPERATIONAL | 30 | 1d04h22m |
| | 4.4.4.4 | cd6/1 | Passive | OPERATIONAL | 30 | 1d04h22m |

```
PE3#show mpls vpls detail
```

```
Virtual Private LAN Service Instance: VPLS101, ID: 101
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:
```

```
Configured interfaces:
```

```
Interface: cd2/1.101
Status: Up
Subinterface Match Criteria(s) :
dot1q 1
```

```
Mesh Peers:
```

```
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d04h20m)
FEC signaling element: FEC128
2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d04h20m)
FEC signaling element: FEC128
```

```
PE3#show mpls vpls VPLS101
```

```
Virtual Private LAN Service Instance: VPLS101, ID: 101
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:
```

```
Configured interfaces:
```

```
Interface: cd2/1.101
Status: Up
Subinterface Match Criteria(s) :
dot1q 1
```

```
Mesh Peers:
```

```
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d04h21m)
FEC signaling element: FEC128
2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d04h21m)
FEC signaling element: FEC128
```



```
PE3#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|---|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | | |
| 101 | 1.1.1.1 | 402 | 24961 | cd6/1 | 24961 | 2/Up | 2 | L |
| DP | Active | 1d04h21m | - | | | | | |
| 101 | 2.2.2.2 | 302 | 24960 | cd6/1 | 24973 | 2/Up | 4 | L |
| DP | Active | 1d04h21m | - | | | | | |

VPLS - Tunnel-ID

Overview


Virtual Private LAN Service (VPLS) is a Layer 2 VPN technology that allows geographically dispersed sites to share an Ethernet broadcast domain by connecting them through a service provider's MPLS network. This effectively creates a single, bridged Ethernet LAN across multiple sites.

VPLS with Tunnel-IDs provides a robust solution for extending Ethernet LAN services over a wide area network. The use of Tunnel-IDs ensures efficient and reliable delivery of Ethernet frames across the MPLS core, enabling service providers to offer high-quality, scalable, and flexible VPN services to their customers.

Feature Characteristics

Tunnel-IDs in VPLS are essential for efficient and reliable operation due to their unique identification, support for scalability, dynamic allocation by protocols like LDP or BGP, and compatibility with various signaling methods. They enable precise traffic engineering and load balancing, enhance resilience through fast reroute mechanisms and redundancy, and improve security with traffic isolation and access control. Additionally, Tunnel-IDs facilitate network management and monitoring by providing visibility into network paths and performance metrics, and they ensure compatibility with existing infrastructure for seamless integration of new services.

The examples show the minimum configuration required for enabling a VPLS Tunnel-ID between PE1, PE2, and PE3.

Figure 70. VPLS Tunnel-ID

Configuration

The VPLS Tunnel-ID feature enables you to use BGP as the control plane protocol for both auto discovery and signaling in accordance with RFC 4761

PE1 Configuration

1. To enable the VPLS-Tunnel-ID, execute the following command in the config mode

```
PE1(config)#configure terminal
PE1(config)#interface lo
PE1(config-if)#ip address 1.1.1.1/32 secondary
PE1(config-if)#exit
PE1(config)#interface ce2/1
PE1(config-if)#ip address 10.1.1.1/24
PE1(config-if)#exit
PE1(config)#router ospf 100
PE1(config-router)#network 10.1.1.0/24 area 0
PE1(config-router)#network 1.1.1.1/32 area 0
PE1(config-router)#exit
PE1(config)#router ldp
PE1(config-router)# router-id 1.1.1.1
PE1(config-router)#transport-address ipv4 1.1.1.1
PE1(config-router)#targeted-peer ipv4 2.2.2.2
PE1(config-router-targeted-peer)#exit
PE1(config-router)#targeted-peer ipv4 3.3.3.3
PE1(config-router-targeted-peer)#exit
PE1(config-router)#exit
PE1(config)#interface ce2/1
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)#exit
PE1(config)# mpls vpls vpls104 104
PE1(config-vpls)#signaling ldp
```

```

PE1(config-vpls-sig)# vpls-peer 2.2.2.2 tunnel-id 701 vpls-type vlan
PE1(config-vpls-sig)# vpls-peer 3.3.3.3 tunnel-id 108 vpls-type ethernet
PE1(config-vpls-sig)#exit-signalig
PE1(config-vpls)#exit
PE1(config)#mpls ftn-entry tunnel-id 107 3.3.3.3/32 523 10.1.1.2 ce2/1 primary
PE1(config)#mpls ftn-entry tunnel-id 701 2.2.2.2/32 522 10.1.1.2 ce2/1 primary
PE1(config)# interface xe2.104 switchport
PE1(config-if)#encapsulation dot1q 4
PE1(config-if)#access-if-vpls
PE1(config-acc-if-vpls)#mpls-vpls vpls104
PE1(config-acc-if-vpls)#commit

```

2. To exit the configuration, execute the following command.

```

PE1(config-if-vpls)#exit
PE1(config-if)#exit

```

3. To save the configuration, execute the following command.

```

PE1(config)#commit
PE1(config)#copy running-config startup-config

```

P Configuration

1. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```

P(config)#configure terminal
P(config)#interface lo
P(config-if)#ip address 4.4.4.4/32 secondary
P(config-if)#exit
P(config)#interface xel1
P(config-if)#ip address 10.10.1.2/24
P(config-if)#exit
P(config)#interface xel3
P(config-if)#ip address 20.1.1.2/24
P(config-if)#exit
P(config)#interface xel2
P(config-if)#ip address 30.1.1.2/24
P(config-if)#exit
P(config)#router ospf 100
P(config-router)#network 10.1.1.0/24 area 0
P(config-router)#network 20.1.1.0/24 area 0
P(config-router)#network 30.1.1.0/24 area 0
P(config-router)#network 4.4.4.4/32 area 0
P(config-router)#exit
P(config)#router ldp
P(config-router)#router-id 4.4.4.4
P(config-router)#transport-address ipv4 4.4.4.4
P(config-router-targeted-peer)#exit
P(config-router)#exit
P(config)#interface xel1
P(config-if)#label-switching
P(config-if)#enable-ldp ipv4
P(config-if)#enable-rsvp
P(config-if)#exit
P(config)#interface xel2
P(config-if)#label-switching
P(config-if)#enable-ldp ipv4
P(config-if)#enable-rsvp
P(config-if)#exit
P(config)#interface xel3
P(config-if)#label-switching
P(config-if)#enable-ldp ipv4
P(config-if)# enable-rsvp

```

2. To exit the configuration, execute the following command.

```

P(config-if)#exit

```

3. To save the configuration, execute the following command.

```
P(config)#commit
P(config)#copy running-config startup-config
```

PE2 Configuration

1. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```
PE2(config)#configure terminal
PE2(config)#interface lo
PE2(config-if)#ip address 2.2.2.2/32 secondary
PE2(config-if)#exit
PE2(config)#interface xe15
PE2(config-if)#ip address 20.1.1.1/24
PE2(config-if)#no shutdown
PE2(config-if)#exit
PE2(config)#router ospf 100
PE2(config-router)#network 20.1.1.0/24 area 0
PE2(config-router)#network 2.2.2.2/32 area 0
PE2(config-router)#exit
PE2(config)#router ldp
PE2(config-router)#router-id 2.2.2.2
PE2(config-router)#transport-address ipv4 2.2.2.2
PE2(config-router)#targeted-peer ipv4 1.1.1.1
PE2(config-router-targeted-peer)#exit
PE2(config-router)#targeted-peer ipv4 3.3.3.3
PE2(config-router-targeted-peer)#exit
PE2(config-router)#exit
PE2(config)#interface xe15
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)#exit
PE2(config)# mpls vpls vpls104 104
PE2(config-vpls)#signaling ldp
PE2(config-vpls-sig)# vpls-peer 1.1.1.1 tunnel-id 701 vpls-type vlan
PE2(config-vpls-sig)# vpls-peer 3.3.3.3 tunnel-id 108 vpls-type ethernet
PE2(config-vpls)#exit
PE2(config-vpls)#commit
PE2(config)# mpls ftn-entry tunnel-id 701 1.1.1.1/32 523 20.1.1.2 xe15 primary
PE2(config)# mpls ftn-entry tunnel-id 108 3.3.3.3/32 5534 20.1.1.2 xe15 primary
PE2(config)#int xe0.104 switchport
PE2(config-if)#encapsulation dot1q 4
PE2(config-if)#access-if-vpls
PE2(config-acc-if-vpls)#mpls-vpls VPLS104
```

2. To exit the configuration, execute the following command.

```
PE2(config-if-vpls)#exit
PE2(config-if)#exit
```

3. To save the configuration, execute the following command.

```
PE2(config)#commit
PE2(config)#copy running-config startup-config
```

PE3 Configuration

1. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```
PE3(config)#configure terminal
PE3(config)#interface lo
PE3(config-if)#ip address 3.3.3.3/32 secondary
PE3(config-if)#exit
PE3(config)#interface cd6/1
PE3(config-if)#ip address 30.1.1.1/24
PE3(config-if)#no shutdown
PE3(config-if)#exit
```

```

PE3(config)#router ospf 100
PE3(config-router)#network 30.1.1.0/24 area 0
PE3(config-router)#network 3.3.3.3/32 area 0
PE3(config-router)#exit
PE3(config)#router ldp
PE3(config-router)#router-id 3.3.3.3
PE3(config-router)#transport-address ipv4 3.3.3.3
PE3(config-router)#targeted-peer ipv4 1.1.1.1
PE3(config-router-targeted-peer)#exit
PE3(config-router)#targeted-peer ipv4 2.2.2.2
PE3(config-router-targeted-peer)#exit
PE3(config-router)#exit
PE3(config)#interface cd6/1
PE3(config-if)#label-switching
PE3(config-if)#enable-ldp ipv4
PE3(config-if)#exit
PE3(config)#mpls vpls VPLS104 104
PE3(config-vpls)#signaling ldp
PE3(config-vpls-sig)#vpls-peer 2.2.2.2 tunnel-id 108 vpls-type Ethernet
PE3(config-vpls-sig)#vpls-peer 1.1.1.1 tunnel-id 701 vpls-type Ethernet
PE3(config-vpls-sig)#exit-signaling
PE3(config-vpls)#exit
PE3(config-vpls-sig)# mpls ftn-entry tunnel-id 108 2.2.2.2/32 524 30.1.1.2 cd6/1 primary
PE3(config-vpls-sig)# mpls ftn-entry tunnel-id 701 1.1.1.1/32 524 30.1.1.2 cd6/1 primary
PE3(config)#int cd2/1.104 switchport
PE3(config-if)#encapsulation dot1q 4
PE3(config-if)#access-if-vpls
PE3(config-acc-if-vpls)#mpls-vpls VPLS104

```

2. To exit the configuration, execute the following command.

```

PE3(config-if-vpls)#exit
PE3(config-if)#exit

```

3. To save the configuration, execute the following command.

```

PE3(config)#commit
PE3(config)#copy running-config startup-config

```

Validation

Use this command to validate the Per peer VPLS-type signaling.

PE1

```

PE1#show mpls vpls detail
Virtual Private LAN Service Instance: vpls104, ID: 104
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe2.104
Status: Up
Subinterface Match Criteria(s) :
dot1q 4

Mesh Peers:
2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:16:55)

```

```

    FEC signaling element: FEC128
    Tunnel-Id: 701
    3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:24:45)
    FEC signaling element: FEC128
    Tunnel-Id: 107

PE1#sh mpls vpls vpls104
Virtual Private LAN Service Instance: vpls104, ID: 104
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe2.104
Status: Up
Subinterface Match Criteria(s) :
dot1q 4

Mesh Peers:
    2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:17:31)
    FEC signaling element: FEC128
    Tunnel-Id: 701
    3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:25:21)
    FEC signaling element: FEC128
    Tunnel-Id: 107

PE1#sh ldp vpls 104

VPLS Identifier      : 104
Peer IP              : 2.2.2.2
VC State             : UP
VC Type              : vlan
VC Label Sent        : 24969
VC Label Received    : 24976
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
Forwarding
Active
Remote PW Status    :
Forwarding
Active
Local FEC Type       : 128
Remote FEC Type      : 128
LDP-VPLS Signaled Time : 00:18:11

VPLS Identifier      : 104
Peer IP              : 3.3.3.3
VC State             : UP
VC Type              : ethernet
VC Label Sent        : 24972
VC Label Received    : 25603

```

```

Local MTU      : 1500
Remote MTU     : 1500
Local Control Word : disabled
Remote Control Word: disabled
Current use    : disabled
Local Flow Label : disabled
Remote Flow Label : disabled
Current use    : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status :
  Forwarding
  Active
Remote PW Status :
  Forwarding
  Active
Local FEC Type : 128
Remote FEC Type : 128
LDP-VPLS Signaled Time : 00:26:01

```

```

PE1#sh mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 104 | 2.2.2.2 | 522 | 24969 | ce2/1 | 24976 | 2/Up | 14 |
| DP | Active | 00:18:51 | - | | | | L |
| 104 | 3.3.3.3 | 523 | 24972 | ce2/1 | 25603 | 2/Up | 16 |
| DP | Active | 00:26:41 | - | | | | L |

```

PE1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|-----------|
| | 2.2.2.2 | ce2/1 | Passive | OPERATIONAL | 30 | 5d23h28m |
| | 3.3.3.3 | ce2/1 | Passive | OPERATIONAL | 30 | 00:27:21 |
| | 4.4.4.4 | ce2/1 | Passive | OPERATIONAL | 30 | 5d23h27m! |

PE2

```
PE2#show mpls vpls VPLS104
```

```

Virtual Private LAN Service Instance: vpls104, ID: 104
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

```

```

Configured interfaces:
Interface: xe0.104
Status: Up
Subinterface Match Criteria(s) :
dot1q 4

```

```

Mesh Peers:
1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:20:44)

```

```

FEC signaling element: FEC128
Tunnel-Id: 701
3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:21:32)
FEC signaling element: FEC128
Tunnel-Id: 108

```

```

FEC signaling element: FEC128

```

```

PE2# show ldp vpls 104
VPLS Identifier      : 104
Peer IP              : 1.1.1.1
VC State              : UP
VC Type              : vlan
VC Label Sent        : 24976
VC Label Received    : 24969
Local MTU             : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
LDP-VPLS Signaled Time : 00:21:22

```

```

VPLS Identifier      : 104
Peer IP              : 3.3.3.3
VC State              : UP
VC Type              : ethernet
VC Label Sent        : 24979
VC Label Received    : 25608
Local MTU             : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
LDP-VPLS Signaled Time : 00:22:10

```

```

PE2#sh mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|---|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | | |
| 104 | 1.1.1.1 | 523 | 24976 | xe15 | 24969 | 2/Up | 17 | L |
| DP | Active | 00:20:31 | - | | | | | |
| 104 | 3.3.3.3 | 534 | 24979 | xe15 | 25608 | 2/Up | 22 | L |


```
DP           Active      00:21:19      -
```

```
PE2#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | xe15 | Active | OPERATIONAL | 30 | 5d23h31m |
| | 3.3.3.3 | xe15 | Passive | OPERATIONAL | 30 | 00:30:26 |
| | 4.4.4.4 | xe15 | Passive | OPERATIONAL | 30 | 01w0d05h |

PE3

```
PE3#show mpls vpls VPLS104
```

```
Virtual Private LAN Service Instance: vpls104, ID: 104
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:
```

```
Configured interfaces:
Interface: cd2/1.104
Status: Up
Subinterface Match Criteria(s) :
dot1q 4
```

```
Mesh Peers:
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:30:03)
FEC signaling element: FEC128
Tunnel-Id: 107
2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:24:09)
FEC signaling element: FEC128
Tunnel-Id: 108
```

```
PE3#sh mpls vpls vpls104
Virtual Private LAN Service Instance: vpls104, ID: 104
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:
```

```
Configured interfaces:
Interface: cd2/1.104
Status: Up
Subinterface Match Criteria(s) :
dot1q 4
```

```
Mesh Peers:
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:30:26)
FEC signaling element: FEC128
Tunnel-Id: 107
2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:24:32)
```

```
FEC signaling element: FEC128
Tunnel-Id: 108
```

```
PE3#sh ldp vpls 104
```

```
VPLS Identifier      : 104
Peer IP              : 1.1.1.1
VC State             : UP
VC Type              : ethernet
VC Label Sent        : 25603
VC Label Received    : 24972
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
Local FEC Type       : 128
Remote FEC Type      : 128
LDP-VPLS Signaled Time : 00:31:45
```

```
VPLS Identifier      : 104
Peer IP              : 2.2.2.2
VC State             : UP
VC Type              : ethernet
VC Label Sent        : 25608
VC Label Received    : 24979
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
Local FEC Type       : 128
Remote FEC Type      : 128
LDP-VPLS Signaled Time : 00:24:49
```

```
PE3#show mpls vpls mesh
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|---|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | | |
| 104 | 1.1.1.1 | 524 | 25603 | cd6/1 | 24972 | 2/Up | 7 | L |
| DP | Active | 00:31:09 | - | | | | | |
| 104 | 2.2.2.2 | 532 | 25608 | cd6/1 | 24979 | 2/Up | 10 | L |
| DP | Active | 00:25:15 | - | | | | | |

VPLS - Tunnel-Name

Overview

Virtual Private LAN Service (VPLS) is a Layer 2 VPN technology that allows geographically dispersed sites to share a common Ethernet broadcast domain, functioning as if they were all connected to the same LAN. This setup enables any-to-any communication between sites, providing the appearance of a single, cohesive LAN across multiple locations.

VPLS - Tunnel-Name is a configuration element within a VPLS network that is used to identify and manage MPLS tunnels. When configuring a VPLS instance on a network device, the Tunnel-Name is assigned to each MPLS tunnel. This can typically be done in the device's CLI under the VPLS configuration mode.


The "VPLS - Tunnel-Name" is a vital element in the configuration and management of VPLS networks. By providing unique identifiers for MPLS tunnels, it enhances the organization, management, and troubleshooting of VPLS instances, ultimately leading to more efficient and reliable network operations.

Feature Characteristics

The VPLS - Tunnel-Name feature in a VPLS network enhances functionality, manageability, and efficiency by providing unique identification for MPLS tunnels, simplifying network management, and improving troubleshooting. It allows granular control over configurations and policies, supports scalability and future network expansions, enhances security through controlled access and traffic isolation, and ensures redundancy and reliability by managing backup tunnels effectively. Additionally, it streamlines operations and reduces configuration errors, making it an essential tool for maintaining robust and reliable VPLS networks.

The examples show the minimum configuration required for enabling a VPLS Tunnel-Name between PE1, PE2, and PE3.

Figure 71. VPLS Tunnel-Name



Configuration

The VPLS Tunnel-Name feature enables you to use BGP as the control plane protocol for both auto discovery and signaling in accordance with RFC 4761

PE1 Configuration

1. To enable the per peer VPLS-Tunnel-Name, execute the following command in the config mode

```
PE1(config)#configure terminal
PE1(config)#interface lo
PE1(config-if)#ip address 1.1.1.1/32 secondary
PE1(config-if)#exit
PE1(config)#interface ce2/1
PE1(config-if)#ip address 10.1.1.1/24
PE1(config-if)#no shutdown
PE1(config-if)#exit
PE1(config)#router rsvp
PE1(config)interface ce2/1
PE1(config-if)# enable-rsvp
PE1(config-if)#exit
PE1(config)#rsvp-trunk pe2 ipv4
PE1(config-trunk)#from 1.1.1.1
PE1(config-trunk)#to 2.2.2.2
PE1(config-trunk)#exit
PE1(config)#rsvp-trunk pe3 ipv4
PE1(config-trunk)# from 1.1.1.1
PE1(config-trunk)# to 3.3.3.3
PE1(config)#exit
PE1(config)#router ospf 100
PE1(config-router)#network 10.1.1.0/24 area 0
PE1(config-router)#network 1.1.1.1/32 area 0
PE1(config-router)#exit
PE1(config)#router ldp
PE1(config-router)# router-id 1.1.1.1
PE1(config-router)#transport-address ipv4 1.1.1.1
PE1(config-router)#targeted-peer ipv4 2.2.2.2
PE1(config-router-targeted-peer)#exit
PE1(config-router)#targeted-peer ipv4 3.3.3.3
PE1(config-router-targeted-peer)#exit
PE1(config-router)#exit
PE1(config)#interface ce2/1
PE1(config-if)#label-switching
PE1(config-if)#enable-ldp ipv4
PE1(config-if)#exit
PE1(config)#mpls vpls vpls111 111
PE1(config-vpls)#signaling ldp
PE1(config-vpls-sig)#vpls-peer 2.2.2.2 tunnel-name pe2 vpls-type vlan
PE1(config-vpls-sig)#vpls-peer 3.3.3.3 tunnel-name pe3 vpls-type ethernet
PE1(config-vpls-sig)#exit-signaling
PE1(config-vpls)#exit
PE1(config)#int xe2.222 switchport
PE1(config-if)#encapsulation dot1q 222
PE1(config-if-vpls)#access-if-vpls
PE1(config-acc-if-vpls)#mpls-vpls VPLS111
PE1(config-acc-if-vpls)#commit
```

2. To exit the configuration, execute the following command.

```
PE1(config-if-vpls)#exit
PE1(config-if)exit
```

3. To save the configuration, execute the following command.

```
PE1(config)#commit
PE1(config)#copy running-config startup-config
```

P Configuration

1. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```
P(config)#configure terminal
P(config)#interface lo
P(config-if)#ip address 4.4.4.4/32 secondary
P(config-if)#exit
P(config)#interface xel1
P(config-if)#ip address 10.1.1.2/24
P(config-if)#exit
P(config)#interface xel3
P(config-if)#ip address 20.2.1.2/24
P(config-if)#exit
P(config)#interface xel2
P(config-if)#ip address 30.3.1.2/24
P(config-if)#exit
P(config)#router ospf 100
P(config-router)#network 10.1.1.0/24 area 0
P(config-router)#network 20.2.1.0/24 area 0
P(config-router)#network 30.3.1.0/24 area 0
P(config-router)#network 4.4.4.4/32 area 0
P(config-router)#exit
P(config)#router ldp
P(config-router)#router-id 4.4.4.4
P(config-router)#transport-address ipv4 4.4.4.4
P(config-router-targeted-peer)#exit
P(config-router)#exit
P(config)#interface xel1
P(config-if)#label-switching
P(config-if)#enable-ldp ipv4
P(config-if)#enable-rsvp
P(config-if)#exit
P(config)#interface xel2
P(config-if)#label-switching
P(config-if)#enable-ldp ipv4
P(config-if)#enable-rsvp
P(config-if)#exit
P(config)#interface xel3
P(config-if)#label-switching
P(config-if)#enable-ldp ipv4
P(config-if)#enable-rsvp
```

2. To exit the configuration, execute the following command.

```
P(config-if)#exit
```

3. To save the configuration, execute the following command.

```
P(config)#commit
P(config)#copy running-config startup-config
```

PE2 Configuration

4. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```
PE2(config)#configure terminal
PE2(config)#interface lo
PE2(config-if)#ip address 2.2.2.2/32 secondary
PE2(config-if)#exit
PE2(config)#interface xel5
PE2(config-if)#ip address 20.1.1.1/24
PE2(config-if)#no shutdown
PE2(config-if)#exit
PE2(config)#router ospf 100
PE2(config-router)#network 20.1.1.0/24 area 0
PE2(config-router)#network 2.2.2.2/32 area 0
PE2(config-router)#exit
```

```

PE2(config)#router rsvp
PE2(config-router)#exit
PE2(config)#interface xe15
PE2(config-if)#enable-rsvp
PE2(config-if)#exit
PE2(config)#rsvp-trunk pe1 ipv4
PE2(config-trunk)#from 2.2.2.2
PE2(config-trunk)#to 1.1.1.1
PE2(config-trunk)#exit
PE2(config)#rsvp-trunk pe3 ipv4
PE2(config-trunk)#from 2.2.2.2
PE2(config-trunk)#to 3.3.3.3
PE2(config-trunk)#exit
PE2(config)#router ldp
PE2(config-router)#router-id 2.2.2.2
PE2(config-router)#transport-address ipv4 2.2.2.2
PE2(config-router)#targeted-peer ipv4 1.1.1.1
PE2(config-router-targeted-peer)#exit
PE2(config-router)#targeted-peer ipv4 3.3.3.3
PE2(config-router-targeted-peer)#exit
PE2(config-router)#exit
PE2(config)#interface xe15
PE2(config-if)#label-switching
PE2(config-if)#enable-ldp ipv4
PE2(config-if)#enable-rsvp
PE2(config-if)#exit
PE2(config)#mpls vpls VPLS111 111
PE2(config-vpls)#signaling ldp
PE2(config-vpls-sig)#vpls-peer 1.1.1.1 tunnel-name pe1 vpls-type vlan
PE2(config-vpls-sig)#vpls-peer 3.3.3.3 tunnel-name pe3 vpls-type ethernet
PE2(config-vpls-sig)#exit-signaling
PE2(config-vpls)#exit
PE2(config)#int xe0.222 switchport
PE2(config-if)#encapsulation dot1q 222
PE2(config-if)#access-if-vpls
PE2(config-acc-if-vpls)#mpls-vpls vpls111
PE2(config-acc-if-vpls)#commit

```

5. To exit the configuration, execute the following command.

```

PE2(config-if-vpls)#exit
PE2(config-if)exit

```

6. To save the configuration, execute the following command.

```

PE2(config)#commit
PE2(config)#copy running-config startup-config

```

PE3 Configuration

1. To enable the per peer VPLS-type signaling, execute the following command in the config mode

```

PE3(config)#configure terminal
PE3(config)#interface lo
PE3(config-if)#ip address 3.3.3.3/32 secondary
PE3(config-if)#exit
PE3(config)#interface cd6/1
PE3(config-if)#ip address 30.1.1.1/24
PE3(config-if)#no shutdown
PE3(config-if)#exit
PE3(config)#router rsvp
PE3(config-router)#exit
PE3(config)#rsvp-trunk pe1_pe3 ipv4
PE3(config-trunk)#from 3.3.3.3
PE3(config-trunk)#to 1.1.1.1
PE3(config-trunk)#rsvp-trunk pe3_pe2 ipv4

```

```

PE3(config-trunk)#from 3.3.3.3
PE3(config-trunk)#to 2.2.2.2
PE3(config-trunk)#commit
PE3(config-trunk)#exit
PE3(config)#router ospf 100
PE3(config-router)#network 30.1.1.0/24 area 0
PE3(config-router)#network 3.3.3.3/32 area 0
PE3(config-router)#exit
PE3(config)#router ldp
PE3(config-router)#router-id 3.3.3.3
PE3(config-router)#transport-address ipv4 3.3.3.3
PE3(config-router)#targeted-peer ipv4 1.1.1.1
PE3(config-router-targeted-peer)#exit
PE3(config-router)#targeted-peer ipv4 2.2.2.2
PE3(config-router-targeted-peer)#exit
PE3(config-router)#exit
PE3(config)#interface cd6/1
PE3(config-if)#label-switching
PE3(config-if)#enable-ldp ipv4
PE3(config-if)#enable-rsvp
PE3(config-if)#exit
PE3(config)#mpls vpls VPLS111 111
PE3(config-vpls)#signaling ldp
PE3(config-vpls-sig)#vpls-peer 2.2.2.2 tunnel-name pe3_pe2 vpls-type vlan
PE3(config-vpls-sig)#vpls-peer 1.1.1.1 tunnel-name pe1_pe3 vpls-type tunnel-name ethernet
PE3(config-vpls-sig)#exit-signaling
PE3(config-vpls)#exit
PE3(config)#interface cd2/1.222 switchport
PE3(config-if)# encapsulation dot1q 222
PE3(config-if)# access-if-vpls
PE3(config-acc-if-vpls)# mpls-vpls VPLS111
PE3(config-acc-if-vpls)#commit

```

2. To exit the configuration, execute the following command.

```

PE3(config-if-vpls)#exit
PE3(config-if)#exit

```

3. To save the configuration, execute the following command.

```

PE3(config)#commit
PE3(config)#copy running-config startup-config

```

Validation

Use this command to validate the Per peer VPLS-type signaling.

PE1 Validation

```

PE1#show mpls vpls detail
Virtual Private LAN Service Instance: vpls111, ID: 111
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:

```

```
Interface: xe2.222
Status: Up
Subinterface Match Criteria(s) :
dot1q 222
```

Mesh Peers:

```
2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:49:42)
FEC signaling element: FEC128
Tunnel-Name: pe2
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:48:36)
FEC signaling element: FEC128
Tunnel-Name: pe3
```

```
PE1#show mpls vpls mesh
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | SIG- |
|----------|-----------|--------------|-----------|--------------|-----------|---------|----------|------|
| Protocol | Status | UpTime | Ext-Color | | | | | |
| 111 | 2.2.2.2 | 25604 | 24971 | ce2/1 | 24977 | 2/Up | 13 | LDP |
| | Active | 00:50:49 | - | | | | | |
| 111 | 3.3.3.3 | 25603 | 24974 | ce2/1 | 25604 | 2/Up | 15 | LDP |
| | Active | 00:49:43 | - | | | | | |

```
PE1#sh rsvp session
```

```
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary
```

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|---------|---------|--------|----------|----------|---------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 2.2.2.2 | 1.1.1.1 | 5001 | 2201 | PRI | pe2- | |
| Primary | | UP | 02:35:04 | 1 1 SE | - | 25604 |
| 3.3.3.3 | 1.1.1.1 | 5002 | 2202 | PRI | pe3- | |
| Primary | | UP | 00:50:16 | 1 1 SE | - | 25603 |

Total 2 displayed, Up 2, Down 0.

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|---------|---------|--------|----------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 1.1.1.1 | 2.2.2.2 | 5001 | 2201 | PRI | 1- | |
| Primary | | UP | 02:35:04 | 1 1 SE | 25600 | - |
| 1.1.1.1 | 3.3.3.3 | 5001 | 2201 | PRI | pe1_pe3- | |
| Primary | | UP | 00:50:12 | 1 1 SE | 25602 | - |
| 1.1.1.1 | 2.2.2.2 | 5002 | 2202 | PRI | pe1- | |
| Primary | | UP | 02:35:04 | 1 1 SE | 25601 | - |

Total 3 displayed, Up 3, Down 0.

```
PE1#show ldp vpls 111
```

```
VPLS Identifier : 111
Peer IP : 2.2.2.2
VC State : UP
VC Type : vlan
VC Label Sent : 24971
VC Label Received : 24977
Local MTU : 1500
Remote MTU : 1500
Local Control Word : disabled
Remote Control Word : disabled
Current use : disabled
Local Flow Label : disabled
Remote Flow Label : disabled
Current use : disabled
Local PW Status Capability : enabled
```



```
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status      :
  Forwarding
  Active
Local FEC Type        : 128
Remote FEC Type        : 128
LDP-VPLS Signaled Time : 04:34:22
```

```
VPLS Identifier      : 111
Peer IP              : 3.3.3.3
VC State             : UP
VC Type              : vlan
VC Label Sent        : 24974
VC Label Received    : 25604
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status      :
  Forwarding
  Active
Local FEC Type        : 128
Remote FEC Type        : 128
LDP-VPLS Signaled Time : 00:50:52
```

```
VPLS Identifier      : 101
Peer IP              : 2.2.2.2
VC State             : UP
VC Type              : vlan
VC Label Sent        : 24960
VC Label Received    : 24972
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status      :
  Forwarding
  Active
Local FEC Type        : 128
Remote FEC Type        : 128
LDP-VPLS Signaled Time : 1d05h00m
```

```
VPLS Identifier      : 101
Peer IP              : 3.3.3.3
```

```

VC State          : UP
VC Type           : ethernet
VC Label Sent     : 24961
VC Label Received : 24961
Local MTU         : 1500
Remote MTU        : 1500
Local Control Word : disabled
Remote Control Word : disabled
Current use       : disabled
Local Flow Label   : disabled
Remote Flow Label  : disabled
Current use       : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status   :
  Forwarding
  Active
Remote PW Status :
  Forwarding
  Active
Local FEC Type    : 128
Remote FEC Type   : 128
LDP-VPLS Signaled Time : 1d04h07m

```

PE1#show mpls vpls mesh

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 111 | 2.2.2.2 | 120 | 24960 | ce2/1 | 24972 | 2/Up | 11 |
| DP | Active | 1d05h01m | - | | | | L |
| 111 | 3.3.3.3 | 103 | 24961 | ce2/1 | 24961 | 2/Up | 12 |
| DP | Active | 1d04h05m | - | | | | L |

PE1#show ldp session

```

Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 2.2.2.2 | ce2/1 | Passive | OPERATIONAL | 30 | 1d05h20m |
| | 3.3.3.3 | ce2/1 | Passive | OPERATIONAL | 30 | 1d04h09m |
| | 4.4.4.4 | ce2/1 | Passive | OPERATIONAL | 30 | 1d05h20m |

PE1#show ip ospf neighbor

Total number of full neighbors: 1

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|----------|-----------|-------------|
| 4.4.4.4 | 1 | Full/Backup | 00:00:35 | 10.1.1.2 | ce2/1 | 0 |

PE2

PE2#sh rsvp session

```

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|---------|---------|--------|----------|----------|---------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 1.1.1.1 | 2.2.2.2 | 5001 | 2201 | PRI | 1- | |
| Primary | | UP | 02:37:30 | 1 1 SE | - | 25600 |
| 1.1.1.1 | 2.2.2.2 | 5002 | 2202 | PRI | pe1- | |
| Primary | | UP | 02:37:30 | 1 1 SE | - | 25601 |
| 3.3.3.3 | 2.2.2.2 | 5003 | 2203 | PRI | pe3- | |

```
Primary          UP    00:52:48  1 1 SE    -        25602
Total 3 displayed, Up 3, Down 0.
```

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|---------|---------|--------|----------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 2.2.2.2 | 1.1.1.1 | 5001 | 2201 | PRI | pe2- | |
| Primary | | UP | 02:37:30 | 1 1 SE | 26240 | - |
| 2.2.2.2 | 3.3.3.3 | 5002 | 2202 | PRI | pe3_pe2- | |
| Primary | | UP | 00:52:38 | 1 1 SE | 26241 | - |

Total 2 displayed, Up 2, Down 0.

```
PE2#show mpls vpls mesh
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 111 | 1.1.1.1 | 25601 | 24977 | xe15 | 24971 | 2/Up | 20 |
| DP | Active | 00:54:36 | - | | | | L |
| 111 | 3.3.3.3 | 25602 | 24978 | xe15 | 25609 | 2/Up | 21 |
| DP | Active | 00:43:11 | - | | | | L |

```
PE2#show mpls vpls VPLS111
```

```
Virtual Private LAN Service Instance: vpls111, ID: 111
```

```
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:
```

Configured interfaces:

```
Interface: xe0.222
Status: Up
Subinterface Match Criteria(s) :
dot1q 222
```

Mesh Peers:

```
1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:55:22)
  FEC signaling element: FEC128
  Tunnel-Name: pe1
3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:43:57)
  FEC signaling element: FEC128
  Tunnel-Name: pe3

  FEC signaling element: FEC128
```

```
PE2#show ldp vpls 111
```

```
VPLS Identifier      : 111
Peer IP              : 1.1.1.1
VC State             : UP
VC Type              : vlan
VC Label Sent        : 24977
VC Label Received    : 24971
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
```

```

Local Flow Label      : disabled
Remote Flow Label     : disabled
Current use           : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status       :
  Forwarding
  Active
Remote PW Status      :
  Forwarding
  Active
LDP-VPLS Signaled Time : 04:38:40

```

```

VPLS Identifier       : 111
Peer IP               : 3.3.3.3
VC State              : UP
VC Type               : ethernet
VC Label Sent         : 24978
VC Label Received     : 25609
Local MTU             : 1500
Remote MTU            : 1500
Local Control Word    : disabled
Remote Control Word   : disabled
Current use           : disabled
Local Flow Label      : disabled
Remote Flow Label     : disabled
Current use           : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status       :
  Forwarding
  Active
Remote PW Status      :
  Forwarding
  Active
LDP-VPLS Signaled Time : 00:44:47

```

PE2#show ldp session

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 1.1.1.1 | xe15 | Active | OPERATIONAL | 30 | 00:41:59 |
| 3.3.3.3 | xe15 | Passive | OPERATIONAL | 30 | 00:30:15 |
| 4.4.4.4 | xe15 | Passive | OPERATIONAL | 30 | 00:42:14 |

PE3

PE3#show RSVP session

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
 State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
 * indicates the session is active with local repair at one or more nodes
 (P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|---------|---------|--------|----------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 1.1.1.1 | 3.3.3.3 | 5001 | 2201 | PRI | pe1_pe3- | |
| Primary | | UP | 00:56:43 | 1 1 SE | - | 25605 |
| 2.2.2.2 | 3.3.3.3 | 5002 | 2202 | PRI | pe3_pe2- | |
| Primary | | UP | 00:56:42 | 1 1 SE | - | 25606 |

Total 2 displayed, Up 2, Down 0.

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|---------|---------|--------|----------|----------|---------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 3.3.3.3 | 1.1.1.1 | 5002 | 2202 | PRI | pe3- | |
| Primary | | UP | 00:56:47 | 1 1 SE | 24321 | - |

```

3.3.3.3      2.2.2.2      5003      2203      PRI      pe3-
Primary      UP      00:56:52  1 1 SE      24320      -
Total 2 displayed, Up 2, Down 0.

```

```

PE3#show ldp session
Codes: m - MD5 password is not set/unset.
      g - GR configuration not set/unset.
      t - TCP MSS not set/unset.
Session has to be cleared manually

```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 1.1.1.1 | cd6/1 | Active | OPERATIONAL | 30 | 1d04h22m |
| | 2.2.2.2 | cd6/1 | Active | OPERATIONAL | 30 | 1d04h22m |
| | 4.4.4.4 | cd6/1 | Passive | OPERATIONAL | 30 | 1d04h22m |

```

PE3#show mpls vpls detail
Virtual Private LAN Service Instance: vpn, ID: 111
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

```

```

Configured interfaces:
Interface: cd4/4.222
Status: Up
Subinterface Match Criteria(s) :
dot1q 222

```

```

Mesh Peers:
1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up)
  FEC signaling element: FEC128
  Tunnel-Name: pe1_pe3-Primary
2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up)
  FEC signaling element: FEC128
  Tunnel-Name: pe3_pe2-Primary

```

```

PE3#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 111 | 1.1.1.1 | 25601 | 24977 | ce6/1 | 24971 | 2/Up | 20 |
| LDP | Active | 00:54:36 | - | | | | |
| 111 | 2.2.2.2 | 25602 | 24978 | ce6/1 | 25609 | 2/Up | 21 |
| LDP | Active | 00:43:11 | - | | | | |

```

PE3#show mpls vpls vpls111
Virtual Private LAN Service Instance: vpn, ID: 111
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0

```

```
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: cd2/1.222
Status: Up
Subinterface Match Criteria(s) :
dot1q 222
Mesh Peers:
  1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d04h21m)
    FEC signaling element: FEC128
  2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 1d04h21m)
    FEC signaling element: FEC128

PE3#show ldp vpls 111
VPLS Identifier      : 111
Peer IP              : 1.1.1.1
VC State             : UP
VC Type              : vlan
VC Label Sent        : 25604
VC Label Received    : 24974
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
Local FEC Type       : 128
Remote FEC Type      : 128
LDP-VPLS Signaled Time : 00:59:17

VPLS Identifier      : 111
Peer IP              : 2.2.2.2
VC State             : UP
VC Type              : ethernet
VC Label Sent        : 25609
VC Label Received    : 24978
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
Local FEC Type       : 128
Remote FEC Type      : 128
LDP-VPLS Signaled Time : 00:49:02
```

```

VPLS Identifier      : 101
Peer IP              : 1.1.1.1
VC State             : UP
VC Type              : ethernet
VC Label Sent        : 24961
VC Label Received    : 24961
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
Local FEC Type       : 128
Remote FEC Type      : 128
LDP-VPLS Signaled Time : 1d04h23m

```

```

VPLS Identifier      : 101
Peer IP              : 2.2.2.2
VC State             : UP
VC Type              : vlan
VC Label Sent        : 24960
VC Label Received    : 24973
Local MTU            : 1500
Remote MTU           : 1500
Local Control Word   : disabled
Remote Control Word  : disabled
Current use          : disabled
Local Flow Label     : disabled
Remote Flow Label    : disabled
Current use          : disabled
Local PW Status Capability : enabled
Remote PW Status Capability : enabled
Current PW Status TLV : enabled
Local PW Status      :
  Forwarding
  Active
Remote PW Status     :
  Forwarding
  Active
Local FEC Type       : 128
Remote FEC Type      : 128
LDP-VPLS Signaled Time : 1d04h23m

```

PE3#sh mpls vpls mesh

(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 111 | 1.1.1.1 | 27602 | 25604 | N/A | 24974 | 0/Dn | 8 |
| LDP | Active | 00:54:36 | - | | | | |
| 111 | 2.2.2.2 | 27601 | 25609 | N/A | 24978 | 0/Dn | 9 |
| LDP | Active | 00:43:11 | - | | | | |

L2VPN FAT SUPPORT

This chapter contains configurations of L2VPN with FAT support.

The Flow-Aware Transport (FAT) of pseudowire (PW) over an MPLS packet-switched network facilitates the load-balancing of traffic across LDP or BGP based signaled pseudowire for Virtual Private LAN Services (VPLS) and LDP-based signaled pseudowire for Virtual Private Wire Service (VPWS).



Note: For interoperability cases, the Flow Label must be configured using the interface-param-tlv option to ensure proper exchange of parameters between peers.

The following CLI options are available for this purpose:


flow-label transmit interface-param-tlv

flow-label receive interface-param-tlv

flow-label both interface-param-tlv

Topology

Figure 72. L2VPN configuration topology



Configuration for VPLS with FAT

PE-1

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |

| | |
|---|---|
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 1.1.1.1 | Configure Router-id |
| (config-router)#targeted-peer ipv4 6.6.6.6 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth2 | Enter interface mode |
| (config-if)#ip address 10.1.1.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID, (100). The Process ID should be a unique positive integer to identifying the routing process. |
| (config-router)#ospf router-id 1.1.1.1 | Configure ospf Router-id |
| (config-router)#network 10.1.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#mpls vpls vpls100 100 | Configuring VPLS instance with name and VPLS ID |
| PE1 (config-vpls)#flow-label both interface-param-tlv | Configure flow label based on requirement (we have three options both, transit and receive) with dynamic and static. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig)#vpls-peer 6.6.6.6 | Configuring VPLS mesh peers |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |
| (config-if)#interface eth2.100 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls vpls100 | Associating the VPLS Instance to the attachment circuit interface. |

| | |
|-----------------------------|--------------------------|
| (config-acc-if-vpls)#commit | Commit the configuration |
| (config-acc-if-vpls)#end | Return to privilege mode |

P1

| | |
|---|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 2.2.2.2 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth1 | Enter interface mode |
| (config-if)#ip address 10.1.1.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po10 | Enter interface mode |
| (config-if)#ip address 10.1.2.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po30 | Enter interface mode |
| (config-if)#ip address 10.1.3.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |

| | |
|--|---|
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#interface eth5 | Enter the Interface mode for eth5 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 3.3.3.3 | Configure ospf Router-id |
| (config-router)#network 10.1.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.2.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.4.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

P2

| | |
|---|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 3.3.3.3 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#interface po10 | Enter interface mode |
| (config-if)#ip address 10.1.2.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |

| | |
|--|---|
| (config-if)#interface eth1 | Enter the Interface mode for eth1 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po20 | Enter interface mode |
| (config-if)#ip address 10.1.4.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 3.3.3.3 | Configure ospf Router-id |
| (config-router)#network 10.1.2.0/24 area 0 | Define the interface on which OSPF runs and associate the |
| area ID (0) with the interface. | |
| (config-router)#network 10.1.4.0/24 area 0 | Define the interface on which OSPF runs and associate the |
| area ID (0) with the interface. | |
| (config-router)#network 3.3.3.3/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

P3

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |

| | |
|--|---|
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 4.4.4.4 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#interface po30 | Enter interface mode |
| (config-if)#ip address 10.1.3.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth1 | Enter the Interface mode for eth1 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po40 | Enter interface mode |
| (config-if)#ip address 10.1.5.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer, identifying the routing process |
| (config-router)#ospf router-id 4.4.4.4 | Configure ospf Router-id |
| (config-router)#network 10.1.3.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.5.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 4.4.4.4/32 area 0 | Define the interface on which OSPF runs and |

| | |
|------------------------|---|
| | associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

P4

| | |
|---|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 5.5.5.5/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 5.5.5.5 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth1 | Enter interface mode |
| (config-if)#ip address 10.1.6.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po20 | Enter interface mode |
| (config-if)#ip address 10.1.4.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po40 | Enter interface mode |
| (config-if)#ip address 10.1.5.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |

| | |
|--|---|
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#interface eth5 | Enter the Interface mode for eth5 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 5.5.5.5 | Configure ospf Router-id |
| (config-router)#network 10.1.4.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.5.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.6.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 5.5.5.5/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

PE-2

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 6.6.6.6/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 6.6.6.6 | Configure Router-id |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)#transport-address ipv4 6.6.6.6 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth2 | Enter interface mode |

| | |
|--|--|
| (config-if)#ip address 10.1.6.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID(100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 6.6.6.6 | Configure ospf Router-id |
| (config-router)#network 10.1.6.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 6.6.6.6/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#mpls vpls vpls100 100 | Configuring VPLS instance with name and VPLS ID |
| PE1(config-vpls)#flow-label both interface-param-tlv | Configure flow label based on requirement (we have three options both, transit and receive) with dynamic and static. |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig)#vpls-peer 1.1.1.1 | Configuring VPLS mesh peers |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |
| (config-if)#interface eth2.100 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls vpls100 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the configuration |
| (config-acc-if-vpls)#end | Return to privilege mode |

Validation

PE1

```

---
PE1#show mpls vpls bgp_vpls100 mesh
(m) - Service mapped over multipath transport

VPLS-ID   Peer Addr      Tunnel-Label  In-Label   Network-Intf  Out-Label  Lkps/St  PW-INDEX
SIG-Protocol  Status
100        6.6.6.6        24961        24963      xe4            24963      2/Up     1         LDP
        Active
PE1#

```



```

PE1#show mpls vpls vpls100 detail
Virtual Private LAN Service Instance: vpls100, ID: 100
SIG-Protocol: LDP
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Enabled, Direction: Both, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Configured interfaces:
  Interface: xe18.100
Subinterface Match Criteria(s) :
  dot1q 100
Mesh Peers:
  6.6.6.6 (Up)

```

```

#ping mpls vpls 200 peer 6.6.6.6/32
Sending 5 MPLS Echos to VPLS Id : 100, timeout is 5 seconds

```

```

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

```

Type 'Ctrl+C' to abort

```

!
!
!
!
!

```

Success Rate is 100.00 percent (5/5)

PE2

```

===
PE2#show mpls vpls vpls200 mesh
(m) - Service mapped over multipath transport

VPLS-ID      Peer Addr      Tunnel-Label  In-Label  Network-Intf  Out-Label  Lkps/St  PW-INDEX
SIG-Protocol  Status
200          1.1.1.1        24961        24963     xe4           24963      2/Up     1         BGP
              Active

PE2#
PE1#show mpls vpls vpls100 detail
Virtual Private LAN Service Instance: vpls100, ID: 100
SIG-Protocol: LDP
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Enabled, Direction: Both, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Configured interfaces:
  Interface: xe18.100
Subinterface Match Criteria(s) :

```

```

dot1q 100
Mesh Peers:
  1.1.1.1 (Up)

#ping mpls vpls 200 peer 1.1.1.1/32
Sending 5 MPLS Echos to VPLS Id : 100, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

!
!
!
!
!

Success Rate is 100.00 percent (5/5)

```

P1

```

==
P1#clear interface counters

P1#show interface counters rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
| Po10      | 1549.80 | 1265260 | 1034.60 | 844649 |
| Po30      | 0.01    | 8        | 0.00    | 8       |
| xe4       | 1034.60 | 844647  | 1549.80 | 1265257 |
| xe10      | 0.00    | 0        | 0.00    | 0       |
| xe11      | 0.00    | 0        | 0.00    | 0       |
| xe12      | 943.60  | 773213  | 474.05  | 384785  |
| xe13      | 0.00    | 4        | 0.00    | 4       |
| xe16      | 606.20  | 492047  | 560.54  | 459865  |
| xe17      | 0.00    | 0        | 0.00    | 0       |
| xe19      | 0.00    | 4        | 0.00    | 4       |
P1#

```

P4

```

==
P4#clear interface counters

P4#show interface counters rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
| Po20      | 1549.80 | 1265260 | 1034.60 | 844649 |
| Po40      | 0.01    | 8        | 0.00    | 8       |
| xe4       | 1034.60 | 844647  | 1549.80 | 1265257 |
| xe10      | 0.00    | 0        | 0.00    | 0       |

```

```

xe11          0.00          0          0.00          0
xe12          943.60        773213        474.05        384785
xe13          0.00          4          0.00          4
xe16          606.20        492047        560.54        459865
xe17          0.00          0          0.00          0
xe19          0.00          4          0.00          4
Pl#

```

Configuration for BGP VPLS with FAT

PE-1

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 1.1.1.1 | Configure Router-id |
| (config-router)#targeted-peer ipv4 6.6.6.6 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth2 | Enter interface mode |
| (config-if)#ip address 10.1.1.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 1.1.1.1 | Configure ospf Router-id |
| (config-router)#network 10.1.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |

| | |
|---|---|
| (config-router)#exit | Exit from router mode |
| (config)#mpls vpls bgp_vpls200 200 | Configuring VPLS instance with name and VPLS ID |
| PE1 (config-vpls)#flow-label both interface-param-tlv | Configure flow label based on requirement (we have three options both, transit and receive) with dynamic and static. |
| (config-vpls)#signaling bgp | Enabling BGP signaling for the VPLS instance |
| (config-vpls-sig)#ve-id 500 | Configure VE ID, which is mandatory for BGP VPLS, otherwise, Signaling does not take place. VE ID should be unique per VPLS instance. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |
| (config-if)#interface eth2.200 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls bgp_vpls200 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the configuration |
| (config-acc-if-vpls)#end | Return to privilege mode |
| #configure terminal | Enter configuration mode. |
| (config)#router bgp 100 | Enter BGP router mode. |
| (config-router)#neighbor 6.6.6.6 remote-as 100 | Configure PE2 as an iBGP peer. |
| (config-router)#neighbor 6.6.6.6 updatesource lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| (config-router)#address-family l2vpn vpls | Configure address-family l2vpn vpls. |
| (config-router-af)#neighbor 6.6.6.6 activate | Activate PE2 in the VPLS address family. |
| (config-router-af)#exit-address-family | Exit address family mode |
| (config-router)#exit | Exit router mode. |

P1

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 2.2.2.2 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |

| | |
|--|---|
| (config-router)#exit | Exit from router mode |
| (config)#interface eth1 | Enter interface mode |
| (config-if)#ip address 10.1.1.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po10 | Enter interface mode |
| (config-if)#ip address 10.1.2.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po30 | Enter interface mode |
| (config-if)#ip address 10.1.3.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#interface eth5 | Enter the Interface mode for eth5 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 3.3.3.3 | Configure ospf Router-id |
| (config-router)#network 10.1.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |

| | |
|--|---|
| (config-router)#network 10.1.2.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.4.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

P2

| | |
|---|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 3.3.3.3 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#interface po10 | Enter interface mode |
| (config-if)#ip address 10.1.2.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth1 | Enter the Interface mode for eth1 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po20 | Enter interface mode |
| (config-if)#ip address 10.1.4.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |

| | |
|--|---|
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 3.3.3.3 | Configure ospf Router-id |
| (config-router)#network 10.1.2.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.4.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 3.3.3.3/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

P3

| | |
|---|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 4.4.4.4 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#interface po30 | Enter interface mode |
| (config-if)#ip address 10.1.3.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth1 | Enter the Interface mode for eth1 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |

| | |
|--|---|
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po40 | Enter interface mode |
| (config-if)#ip address 10.1.5.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 4.4.4.4 | Configure ospf Router-id |
| (config-router)#network 10.1.3.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.5.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 4.4.4.4/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

P4

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 5.5.5.5/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 5.5.5.5 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

| | |
|--|---|
| (config)#interface eth1 | Enter interface mode |
| (config-if)#ip address 10.1.6.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po20 | Enter interface mode |
| (config-if)#ip address 10.1.4.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po40 | Enter interface mode |
| (config-if)#ip address 10.1.5.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#interface eth5 | Enter the Interface mode for eth5 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 5.5.5.5 | Configure ospf Router-id |
| (config-router)#network 10.1.4.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.5.0/24 area 0 | Define the interface on which OSPF runs and |

| | |
|--|---|
| | associate the area ID (0) with the interface. |
| (config-router)#network 10.1.6.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 5.5.5.5/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

PE-2

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 6.6.6.6/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 6.6.6.6 | Configure Router-id |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)#transport-address ipv4 6.6.6.6 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth2 | Enter interface mode |
| (config-if)#ip address 10.1.6.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 6.6.6.6 | Configure ospf Router-id |
| (config-router)#network 10.1.6.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 6.6.6.6/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |

| | |
|--|---|
| (config)#mpls vpls bgp_vpls200 200 | Configuring VPLS instance with name and VPLS ID |
| PE1(config-vpls)#flow-label both interface-param-tlv | Configure flow label based on requirement (we have three options both, transit and receive) with dynamic and static. |
| (config-vpls)#signaling bgp | Enabling BGP signaling for the VPLS instance |
| (config-vpls-sig)#ve-id 600 | Configure VE ID, which is mandatory for BGP VPLS, otherwise, Signaling does not take place. VE ID should be unique per VPLS instance. |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |
| (config-if)#interface eth2.200 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 200 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls bgp_vpls200 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the configuration |
| (config-acc-if-vpls)#end | Return to privilege mode |
| (config)#router bgp 100 | Enter BGP router mode. |
| (config-router)#neighbor 1.1.1.1 remote-as 100 | Configure PE1 as an iBGP peer. |
| (config-router)#neighbor 1.1.1.1 updatesource lo | Update the source as loopback for iBGP peering with the remote PE1 router. |
| (config-router)#address-family l2vpn vpls | Configure address-family l2vpn vpls. |
| (config-router-af)#neighbor 1.1.1.1 activate | Activate PE1 in the VPLS address family. |
| (config-router-af)#exit-address-family | Exit address family mode |
| (config-router)#exit | Exit router mode. |

Validation

PE1

```

---
PE1#show mpls vpls bgp_vpls200 mesh
(m) - Service mapped over multipath transport

VPLS-ID      Peer Addr      Tunnel-Label  In-Label      Network-Intf   Out-Label    Lkps/St      PW-INDEX
SIG-Protocol  Status
200          6.6.6.6        24961        24963         xe4            24963        2/Up         1          BGP
              Active

PE1#
PE1#show mpls vpls bgp_vpls200 detail
Virtual Private LAN Service Instance: bgp_vpls200, ID: 200
SIG-Protocol: BGP
Route-Distinguisher :100:200
Route-Target :100:200
VE-ID :500
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled

```

```

Flow Label Status: Enabled, Direction: Receive, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Configured interfaces:
  Interface: eth1.200
Subinterface Match Criteria(s) :
  dot1q 300
Mesh Peers:
  6.6.6.6 (Up)

```

```

#ping mpls bgp_vpls 2 00 peer 6.6.6.6/32
Sending 5 MPLS Echos to VPLS Id : 100, timeout is 5 seconds

```

```

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

```

```
Type 'Ctrl+C' to abort
```

```

!
!
!
!
!

```

```
Success Rate is 100.00 percent (5/5)
```

PE2

```
===
```

```

PE2#show mpls vpls bgp_vpls200 mesh
(m) - Service mapped over multipath transport

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | | | | | | |
| 200 | 1.1.1.1 | 24961 | 24963 | xe4 | 24963 | 2/Up | 1 BGP |
| | Active | | | | | | |

```
PE2#
```

```

PE1#show mpls vpls bgp_vpls200 detail
Virtual Private LAN Service Instance: bgp_vpls200, ID: 200
SIG-Protocol: BGP
  Route-Distinguisher :100:200
  Route-Target :100:200
  VE-ID :500
Attachment-Circuit :UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Enabled, Direction: Transmit, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Configured interfaces:
  Interface: eth1.200
Subinterface Match Criteria(s) :
  dot1q 200
Mesh Peers:
  1.1.1.1 (Up)

```

```
#ping mpls bgp_vpls 200 peer 1.1.1.1/32
Sending 5 MPLS Echos to VPLS Id : 100, timeout is 5 seconds
```

Codes:

```
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed
```

Type 'Ctrl+C' to abort

```
!
!
!
!
!
```

Success Rate is 100.00 percent (5/5)

P1

```
==
```

```
P1#clear interface counters
```

```
P1#show interface counters rate mbps
```

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-----------|---------|---------|---------|---------|
| Po10 | 1549.80 | 1265260 | 1034.60 | 844649 |
| Po30 | 0.01 | 8 | 0.00 | 8 |
| xe4 | 1034.60 | 844647 | 1549.80 | 1265257 |
| xe10 | 0.00 | 0 | 0.00 | 0 |
| xe11 | 0.00 | 0 | 0.00 | 0 |
| xe12 | 943.60 | 773213 | 474.05 | 384785 |
| xe13 | 0.00 | 4 | 0.00 | 4 |
| xe16 | 606.20 | 492047 | 560.54 | 459865 |
| xe17 | 0.00 | 0 | 0.00 | 0 |
| xe19 | 0.00 | 4 | 0.00 | 4 |

```
P1#
```

```
P4
```

```
==
```

```
P4#clear interface counters
```

```
P4#show interface counters rate mbps
```

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-----------|---------|---------|---------|---------|
| Po20 | 1549.80 | 1265260 | 1034.60 | 844649 |
| Po40 | 0.01 | 8 | 0.00 | 8 |
| xe4 | 1034.60 | 844647 | 1549.80 | 1265257 |
| xe10 | 0.00 | 0 | 0.00 | 0 |
| xe11 | 0.00 | 0 | 0.00 | 0 |
| xe12 | 943.60 | 773213 | 474.05 | 384785 |
| xe13 | 0.00 | 4 | 0.00 | 4 |
| xe16 | 606.20 | 492047 | 560.54 | 459865 |
| xe17 | 0.00 | 0 | 0.00 | 0 |
| xe19 | 0.00 | 4 | 0.00 | 4 |

```
P1#
```

Configuration for VPWS with FAT

PE-1

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 1.1.1.1 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config-router)#targeted-peer ipv4 6.6.6.6 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth2 | Enter interface mode |
| (config-if)#ip address 10.1.1.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 1.1.1.1 | Configure ospf Router-id |
| (config-router)#network 10.1.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#mpls l2-circuit VPWS4 400 3.3.3.3 | Configure the VC for PE-2 |
| PE1(config-vpls)#flow-label both interface-param-tlv | Configure flow label based on requirement (we have three options both, transit and receive) with dynamic and static. |

| | |
|--|--|
| (config-pseudowire)#commit | Commit the configuration |
| (config-pseudowire)#exit | Exit from pseudowire configuration mode |
| (config-if)#interface eth2.300 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpws | Access VPWS under sub interface |
| (config-acc-if-vpws)#mpls-l2-circuit VPWS4 primary | Associating the VPWS Instance to the attachment circuit interface. |
| (config-acc-if-vpws)#commit | Commit the configuration |
| (config-acc-if-vpws)#end | Return to privilege mode |

P1

| | |
|---|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 2.2.2.2 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth1 | Enter interface mode |
| (config-if)#ip address 10.1.1.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po10 | Enter interface mode |
| (config-if)#ip address 10.1.2.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#exit | Exit interface mode |

| | |
|--|---|
| (config)#commit | Commit the transaction. |
| (config)#interface po30 | Enter interface mode |
| (config-if)#ip address 10.1.3.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#interface eth5 | Enter the Interface mode for eth5 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 3.3.3.3 | Configure ospf Router-id |
| (config-router)#network 10.1.1.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.2.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.4.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

P2

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 3.3.3.3 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

| | |
|--|---|
| (config)#interface po10 | Enter interface mode |
| (config-if)#ip address 10.1.2.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth1 | Enter the Interface mode for eth1 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 10 mode active | Moving interface to Dynamic LAG 10 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po20 | Enter interface mode |
| (config-if)#ip address 10.1.4.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 3.3.3.3 | Configure ospf Router-id |
| (config-router)#network 10.1.2.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.4.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 3.3.3.3/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

P3

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 4.4.4.4 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#interface po30 | Enter interface mode |
| (config-if)#ip address 10.1.3.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth1 | Enter the Interface mode for eth1 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 30 mode active | Moving interface to Dynamic LAG 30 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#interface po40 | Enter interface mode |
| (config-if)#ip address 10.1.5.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |

| | |
|--|---|
| (config-router)#ospf router-id 4.4.4.4 | Configure ospf Router-id |
| (config-router)#network 10.1.3.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.5.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 4.4.4.4/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

P4

| | |
|---|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 5.5.5.5/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 5.5.5.5 | Configure Router-id |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth1 | Enter interface mode |
| (config-if)#ip address 10.1.6.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface po20 | Enter interface mode |
| (config-if)#ip address 10.1.4.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth2 | Enter the Interface mode for eth2 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#interface eth3 | Enter the Interface mode for eth3 |
| (config-if)#channel-group 20 mode active | Moving interface to Dynamic LAG 20 |
| (config-if)#exit | Exit interface mode |

| | |
|--|---|
| (config)#commit | Commit the transaction. |
| (config)#interface po40 | Enter interface mode |
| (config-if)#ip address 10.1.5.1/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config-if)#interface eth4 | Enter the Interface mode for eth4 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#interface eth5 | Enter the Interface mode for eth5 |
| (config-if)#channel-group 40 mode active | Moving interface to Dynamic LAG 40 |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 5.5.5.5 | Configure ospf Router-id |
| (config-router)#network 10.1.4.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.5.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 10.1.6.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 5.5.5.5/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the transaction. |
| (config-router)#exit | Exit from router mode |

PE-2

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 6.6.6.6/32 secondary | Configure IP address for the loopback interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 6.6.6.6 | Configure Router-id |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |

| | |
|--|---|
| (config-router)#transport-address ipv4 6.6.6.6 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface eth2 | Enter interface mode |
| (config-if)#ip address 10.1.6.2/24 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf 100 | Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process. |
| (config-router)#ospf router-id 6.6.6.6 | Configure ospf Router-id |
| (config-router)#network 10.1.6.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 6.6.6.6/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#mpls l2-circuit VPWS4 400 1.1.1.1 | Configure the VC for PE-2 |
| PE1(config-vpls)#flow-label both interface-param-tlv | Configure flow label based on requirement (we have three options both, transit and receive) with dynamic and static. |
| (config-pseudowire)#commit | Commit the configuration |
| (config-pseudowire)#exit | Exit from pseudowire configuration mode |
| (config-if)#interface eth2.300 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 300 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpws | Access VPWS under sub interface |
| (config-acc-if-vpws)#mpls-l2-circuit VPWS4 primary | Associating the VPWS Instance to the attachment circuit interface. |
| (config-acc-if-vpws)#commit | Commit the configuration |
| (config-acc-if-vpws)#end | Return to privilege mode |

Validation

PE1

```

---
PE1#show mpls vc-table
(m) - Service mapped over multipath transport

VC-ID          Vlan-ID  Inner-Vlan-ID  Access-Intf  Network-Intf  Out Label  Tunnel-

```

```

Label  Nexthop      Status
400    N/A          N/A          xe18.300      xe4          24967      24961      6.6.6.6
Active

```

```

PE1#
PE1#show mpls l2-circuit detail
MPLS Layer-2 Virtual Circuit: VPWS4, id: 300 PW-INDEX: 4 service-tpid: dot1.q
Endpoint: 6.6.6.6
Control Word: 0
Flow Label Status: Enabled, Direction: Both, Static: Yes
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe18.200
Subinterface Match Criteria(s) :
  dot1q 300
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active

```

PE2

```

===
PE2#show mpls vc-table
(m) - Service mapped over multipath transport

VC-ID      Vlan-ID  Inner-Vlan-ID  Access-Intf  Network-Intf  Out Label  Tunnel-
Label  Nexthop      Status
400      N/A        N/A          xe4.300      xe18          24967      24961      1.1.1.1
Active

PE2#
PE2#show mpls l2-circuit detail
MPLS Layer-2 Virtual Circuit: VPWS4, id: 300 PW-INDEX: 4 service-tpid: dot1.q
Endpoint: 1.1.1.1
Control Word: 0
Flow Label Status: Enabled, Direction: Both, Static: Yes
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe4.200
Subinterface Match Criteria(s) :
  dot1q 200
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active

```

P1

```

==
P1#clear interface counters

P1#show interface counters rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
| Po10      | 1549.80 | 1265260 | 1034.60 | 844649 |
| Po30      | 0.01    | 8       | 0.00    | 8       |
| xe4       | 1034.60 | 844647  | 1549.80 | 1265257 |
| xe10      | 0.00    | 0       | 0.00    | 0       |
| xe11      | 0.00    | 0       | 0.00    | 0       |
| xe12      | 943.60  | 773213  | 474.05  | 384785  |
| xe13      | 0.00    | 4       | 0.00    | 4       |
| xe16      | 606.20  | 492047  | 560.54  | 459865  |
| xe17      | 0.00    | 0       | 0.00    | 0       |
| xe19      | 0.00    | 4       | 0.00    | 4       |

```

P1#

P4

```
==
P4#clear interface counters

P4#show interface counters rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
Po20      1549.80 1265260 1034.60 844649
Po40      0.01    8       0.00    8
xe4       1034.60 844647 1549.80 1265257
xe10      0.00    0       0.00    0
xe11      0.00    0       0.00    0
xe12      943.60 773213 474.05 384785
xe13      0.00    4       0.00    4
xe16      606.20 492047 560.54 459865
xe17      0.00    0       0.00    0
xe19      0.00    4       0.00    4
P1#
```

L2VPN Sub Interface Configuration

L2VPN is a layer 2 service where different locations (customer sites) of an enterprise interconnect to form a big LAN segment. All the locations can exchange layer 2 data with each other via this Virtual LAN in a private and secured way.

The MPLS Core-based L2VPN model has two broad divisions:


- Virtual Private Wire Service (VPWS)
- Virtual Private LAN Service (VPLS)

Virtual Private Wire Service or VPWS is a Point-to-Point (P2P) service implementation of L2VPN. It provides layer 2 data flow of the same or different types (FR, ATM, etc.) of L2 services over the MPLS core attaching the two customer sites.

Virtual Private LAN Service (VPLS) is a Point-to-Multipoint (P2MP) and Multipoint-to-Multipoint (MP2MP) L2VPN service. VPLS is designed for applications that require multipoint access across geographically distributed locations. Using VPLS, several customer sites (or distributed Ethernet LANs) can be interconnected to work as a single bridged domain over the MPLS network. In simple terms, VPLS uses the Layer 2 architecture to create multipoint VPNs that connect several sites over a Wide Area Network (WAN) or Metropolitan Area Network (MAN). The different customer sites are connected via the Service Provider's MPLS core network.

Topology

Figure 73. Link detection debounce timer topology



VPLS Configuration

All configuration commands in the table below should be followed for each Router.

PE1

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 |
| (config-router)#net 49.3600.3600.0001.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface |

| | |
|---|---|
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter interface mode |
| (config-if)#ip address 10.1.1.1/30 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 1.1.1.1 | Configure Router-id |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#mpls vpls vpls100 100 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig)#vpls-peer 3.3.3.3 | Configuring VPLS mesh peers |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |
| (config-if)#interface xe20.100 switchport | Enter sub interface mode |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls vpls100 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the configuration |
| (config-acc-if-vpls)#end | Return to privilege mode |

P

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 |
| (config-router)#net 49.3600.3600.0003.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID |

| | |
|---|---|
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface lo | Enter interface mode |
| (config-if)# ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 2.2.2.2 | Configure Router-id |
| (config-router)# transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit router mode |
| (config)#interface xe5 | Enter interface mode |
| (config-if)#ip address 10.1.1.2/30 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe1 | Enter interface mode |
| (config-if)# ip address 11.1.1.2/30 | Configure IP address on interface |
| (config-if)# label-switching | Enable label switching capability on the interface |
| (config-if)# ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)# enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |

PE2

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 |
| (config-router)#net 49.3600.3600.0002.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface lo | Enter interface mode |

| | |
|---|---|
| (config-if)# ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)# router-id 3.3.3.3 | Configure Router-id |
| (config-router)# targeted-peer ipv4 1.1.1.1 | Configuring targeted LDP sessions to PE-1 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)# transport-address ipv4 3.3.3.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface xe1 | Enter interface mode |
| (config-if)#ip address 11.1.1.1/30 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configurations |
| (config-if)#exit | Exit interface mode |
| (config)#mpls vpls vpls100 100 | Configuring VPLS instance with name and VPLS ID |
| (config-vpls)#signaling ldp | Enabling LDP signaling for the VPLS instance |
| (config-vpls-sig)#vpls-peer 1.1.1.1 | Configuring VPLS mesh peers |
| (config-vpls-sig)#exit-signaling | Exit from VPLS signaling mode |
| (config-vpls)#exit-vpls | Exit from VPLS mode |
| (config)#interface xe18.100 switchport | Enter sub interface mode |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if)#encapsulation dot1q 100 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)#mpls-vpls vpls100 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the configuration |
| (config-acc-if-vpls)#end | Return to privilege mode |

Validation

```
#ping mpls vpls 100 peer 3.3.3.3/32
Sending 5 MPLS Echos to VPLS Id : 100, timeout is 5 seconds
```

```
Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
```

```
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed
```

```
Type 'Ctrl+C' to abort
```

```
!
!
!
!
!
```

```
Success Rate is 100.00 percent (5/5)
```

```
#show mpls vpls mesh
```

```
(m) - Service mapped over multipath transport
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | SIG- | |
|----------|-----------|--------------|----------|--------------|-----------|---------|----------|------|---|
| Protocol | Status | | | | | | | | |
| 100 | 3.3.3.3 | 24320 | 24320 | xe5 | 24320 | 2/Up | 1 | LDP | A |
| ctive | | | | | | | | | |

```
#show mpls vpls detail
```

```
Virtual Private LAN Service Instance: vpls100, ID: 100
```

```
SIG-Protocol: LDP
```

```
Attachment-Circuit :UP
```

```
Learning: Enabled
```

```
Control-Word: Disabled
```

```
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
```

```
Description: none
```

```
service-tpid: dot1q
```

```
Operating mode: Raw
```

```
Configured interfaces:
```

```
Interface: xe20.100
```

```
Subinterface Match Criteria(s) :
```

```
dot1q 100
```

```
Mesh Peers:
```

```
3.3.3.3 (Up)
```

VPWS Configuration

All configuration commands in the table below should be followed for each Router.

PE1

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 |
| (config-router)#net 49.3600.3600.0001.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |

| | |
|---|---|
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe5 | Enter interface mode |
| (config-if)#ip address 10.1.1.1/30 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 1.1.1.1 | Configure Router-id |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configuring targeted LDP sessions to PE-2 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#mpls l2-circuit VPWS4 400 3.3.3.3 | Configure the VC for PE-2 |
| (config-pseudowire)#commit | Commit the configuration |
| (config-pseudowire)#exit | Exit from pseudowire configuration mode |
| (config)#interface xe20.400 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 400 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpws | Access VPWS under sub interface |
| (config-acc-if-vpws)#mpls-l2-circuit VPWS4 primary | Associating the VPWS Instance to the attachment circuit interface. |
| (config-acc-if-vpws)#commit | Commit the configuration |
| (config-acc-if-vpws)#end | Return to privilege mode |

P

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 |
| (config-router)#net 49.3600.3600.0003.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface lo | Enter interface mode |
| (config-if)# ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface |

| | |
|---|---|
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#router-id 2.2.2.2 | Configure Router-id |
| (config-router)# transport-address ipv4 2.2.2.2 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit router mode |
| (config)#interface xe5 | Enter interface mode |
| (config-if)#ip address 10.1.1.2/30 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe1 | Enter interface mode |
| (config-if)# ip address 11.1.1.2/30 | Configure IP address on interface |
| (config-if)# label-switching | Enable label switching capability on the interface |
| (config-if)# ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)# enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |

PE2

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#router isis 1 | Create an IS-IS routing instance for area 49 |
| (config-router)#net 49.3600.3600.0002.00 | Set a Network Entity Title for this instance, specifying the area address and the system ID |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface lo | Enter interface mode |
| (config-if)# ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#commit | Commit the configuration |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP |

| | |
|---|---|
| (config-router)# router-id 3.3.3.3 | Configure Router-id |
| (config-router)# targeted-peer ipv4 1.1.1.1 | Configuring targeted LDP sessions to PE-1 |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit from targeted-peer mode |
| (config-router)# transport-address ipv4 3.3.3.3 | Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface |
| (config-router)#commit | Commit the configuration |
| (config-router)#exit | Exit from router mode |
| (config)#interface xe1 | Enter interface mode |
| (config-if)#ip address 11.1.1.1/30 | Configure IP address on interface |
| (config-if)#label-switching | Enable label switching capability on the interface |
| (config-if)#ip router isis 1 | Enable IS-IS routing on an interface for area 49 |
| (config-if)#enable-ldp ipv4 | Enabling LDP on the interface |
| (config-if)#commit | Commit the configurations |
| (config-if)#exit | Exit interface mode |
| (config)#mpls l2-circuit VPWS4 400 1.1.1.1 | Configure the VC for PE-1 |
| (config-pseudowire)#commit | Commit the configuration |
| (config-pseudowire)#exit | Exit from pseudowire configuration mode |
| (config)#interface xe18.400 switchport | Enter sub interface mode |
| (config-if)#encapsulation dot1q 400 | Configure encapsulation under a subinterface |
| (config-if)#access-if-vpws | Access VPWS under sub interface |
| (config-acc-if-vpws)#mpls-l2-circuit VPWS4 primary | Associating the VPWS Instance to the attachment circuit interface. |
| (config-acc-if-vpws)#commit | Commit the configuration |
| (config-acc-if-vpws)#end | Return to privilege mode |
| (config)#mpls l2-circuit VPWS4 400 1.1.1.1 | Configure the VC for PE-1 |
| (config-pseudowire)#commit | Commit the configuration |

Validation

PE1

```
#show mpls l2-circuit
MPLS Layer-2 Virtual Circuit: VPWS4, id: 400 PW-INDEX: 4 Endpoint: 3.3.3.3
Control Word: 0
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe20.400
Subinterface Match Criteria(s) :
  dot1q 400
Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Active
Virtual Circuit is active

#show ldp mpls-l2-circuit
```

| Transport | Client | VC | VC | Local | Remote | Destination |
|-----------|----------|-------|---------------|----------|----------|-------------|
| VC ID | Binding | State | Type | VC Label | VC Label | Address |
| 400 | xe20.400 | UP | Ethernet VLAN | 24322 | 24322 | 3.3.3.3 |

```
#ping mpls l2-circuit 400
```

```
Sending 5 MPLS Echos to VC Id : 400, timeout is 5 seconds
```

```
Codes:
```

```
'!' - Success, 'Q' - request not sent, '.' - timeout,  
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,  
'N' - LBL Mapping Err, 'D' - DS Mismatch,  
'U' - Unknown Interface, 'R' - Transit (LBL Switched),  
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,  
'P' - Protocol Error, 'X' - Unknown code,  
'Z' - Reverse FEC Validation Failed
```

```
Type 'Ctrl+C' to abort
```

```
!  
!  
!  
!  
!
```

```
Success Rate is 100.00 percent (5/5)
```


Static VPLS Configuration

This chapter includes step-by-step configurations for Static VPLS. It also contains an overview of the concepts of Static VPLS.

Overview

Virtual Private LAN Service (VPLS) is a way to provide Ethernet-based multipoint-to-multipoint communication over IP- MPLS networks. It allows geographically-dispersed sites to share an Ethernet broadcast domain by connecting sites through pseudowires. A set of Martini circuits is grouped by a common VPLS identifier to achieve this service objective.

A pseudowire (PW) consists of a pair of point-to-point, single-hop unidirectional LSPs in opposite directions, each identified by a PW label, also called a Virtual Connection (VC) label.


The Label Distribution Protocol (LDP) is used to signal constituent VCs, and the service provider may use either LDP or RSVP-TE or add static provisioning to set up LSP tunnels to transport data through virtual circuits.

The VPLS identifier is exchanged with the labels, so that both PWs can be linked and be associated with a particular VPLS instance.

Configure Static VPLS

In the following examples, VPLS (v1) is configured on PE-2 with Static VPLS-Peers PE-1 and PE-3 using static LSPs.

Figure 74. Static Virtual Private LAN Service Topology



PE-1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#mpls ftn-entry tunnel-id 11 2.2.2.2/32 120 10.1.1.2 ce2/1 primary | Configure MPLS FTN entry for the creation of a static LSP to PE-2. |
| (config)#mpls ftn-entry tunnel-id 22 3.3.3.3/32 103 10.1.1.2 ce2/1 primary | Configure MPLS FTN entry for the creation of a static |

| | |
|--|---|
| | LSP to PE-3. |
| (config)#mpls ilm-entry 201 pop | Configure MPLS ILM entry for the creation of a static LSP to PE-2. |
| (config)#mpls ilm-entry 301 pop | Configure MPLS ILM entry for the creation of a static LSP to PE-3. |
| (config)#mpls vpls v1 100 | Configure VPLS v1 with ID 100 on PE-1. |
| (config-vpls)#vpls-peer 2.2.2.2 tunnel-id 11 manual | Configure PE-2 as a manual VPLS peer using the static LSP tunnel ID 11 |
| (config-vpls)#vpls-peer 3.3.3.3 tunnel-id 22 manual | Configure PE-3 as a manual VPLS peer using the static LSP tunnel ID 22. |
| (config-vpls)#exit | Exit Configure VPLS mode. |
| (config)#interface ce2/1 | Enter interface mode. |
| (config-if)#ip address 10.10.1.1/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)# router-id 1.1.1.1 | Configure router id same as loopback IP address |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure transport address as loopback address. |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 10.10.1.0/24 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config-router)#network 1.1.1.1/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config-router)#exit | Exit router mode |
| (config)#service-template st1 | Template configuration |
| (config-svc)#exit | Exit service template mode |
| (config)#interface xe2 | Enter interface mode. |
| (config)#vpls fib-entry 100 peer 2.2.2.2 1000 xe1 2000 | Configure VPLS FIB entry for VPLS peer PE-2. |

| | |
|--|--|
| (config)#vpls fib-entry 100 peer 3.3.3.3 3000 xe1 4000 | Configure VPLS FIB entry for VPLS peer PE-3. |
| (config)#commit | Commit the transaction. |

P

| | |
|--|---|
| #configure terminal | Enter configure mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe11 | Specify the interface to be configured. |
| (config-if)#ip address 10.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe13 | Enter interface mode |
| (config-if)#ip address 20.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe12 | Enter interface mode |
| (config-if)#ip address 30.1.1.2/24 | Configure IP address for the interface. |
| (config-if)#label-switching | Enable label switching capability on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#ospf router-id 4.4.4.4 | Configure OSPF router id same as loopback IP address |
| (config-router)#transport-address ipv4 4.4.4.4 | Configure transport address as loopback address. |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 10.1.1.0/24 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the Interface address. |
| (config-router)#network 20.1.1.0/24 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the Interface address. |
| (config-router)#network 30.1.1.0/24 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the Interface address. |
| (config-router)#network 4.4.4.4/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the Interface address. |
| (config)#commit | Commit the transaction. |

PE-2

| | |
|---|---|
| #configure terminal | Enter Configure mode |
| (config)#mpls ftn-entry tunnel-id 11 1.1.1.1/32 201 20.1.1.2 xe15 primary | Configure MPLS FTN entry for the creation of a static LSP to PE-1, and designate xe15 as primary. |
| (config)#mpls ftn-entry tunnel-id 33 3.3.3.3/32 301 20.1.1.2 xe15 primary | Configure MPLS FTN entry for the creation of a static LSP to PE-3, and designate xe15 as primary. |
| (config)#mpls ilm-entry 102 pop | Configure MPLS ILM entry for the creation of a static LSP to PE-1. |
| (config)#mpls ilm-entry 302 pop | Configure MPLS ILM entry for the creation of a static LSP to PE-3 |
| (config)#mpls vpls v1 100 | Configure VPLS v1 with ID 100 on PE-2. |
| (config-vpls)#vpls-peer 1.1.1.1 tunnel-id 11 manual | Configure PE-1 as a manual VPLS peer using static LSP tunnel ID |
| (config-vpls)#vpls-peer 3.3.3.3 tunnel-id 33 manual | Configure PE-3 as a manual VPLS peer using static LSP tunnel ID |
| (config-vpls)#exit | Exit Configure VPLS mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe15 | Enter interface mode |
| (config-if)#ip address 20.1.1.1/24 | Configure IP address for the interface |
| (config-if)#label-switching | Configure label switching |
| (config-if)#exit | Exit interface mode |
| (config)#vpls fib-entry 100 peer 1.1.1.1 2500 xe15 1500 | Configure VPLS FIB entry for VPLS peer PE-1. |
| (config)#vpls fib-entry 100 peer 3.3.3.3 3500 xe15 4500 | Configure VPLS FIB entry for VPLS peer PE-3. |
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 2.2.2.2 | Configure router ID same as loopback IP address |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure transport address as loopback address. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 20.1.1.0/24 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |

| | |
|---|---|
| (config-router)#network 2.2.2.2/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config-router)#exit | Exit router mode |
| (config)#commit | Commit the transaction. |

PE-3

| | |
|---|---|
| #configure terminal | Enter Configure mode |
| (config)#mpls ftn-entry tunnel-id 11 1.1.1.1/32 301 30.30.1.2 cd6/1 primary | Configure MPLS FTN entry for the creation of a static LSP to PE-1. |
| (config)#mpls ftn-entry tunnel-id 22 2.2.2.2/32 302 30.30.1.2 cd6/1 primary | Configure MPLS FTN entry for the creation of a static LSP to PE-2. |
| (config)#mpls ilm-entry 103 pop | Configure MPLS ILM entry for the creation of a static LSP to PE-1. |
| (config)#mpls ilm-entry 203 pop | Configure MPLS ILM entry for the creation of a static LSP to PE-2. |
| (config)#mpls vpls v1 100 | Configure VPLS v1 with ID 100 on PE-3. |
| (config-vpls)#vpls-peer 1.1.1.1 tunnel-id 11 manual | Configure PE-1 as a manual VPLS peer using static LSP tunnel ID 11. |
| (config-vpls)#vpls-peer 2.2.2.2 tunnel-id 22 manual | Configure PE-2 as a manual VPLS peer using static LSP tunnel ID 22. |
| (config-vpls)#exit | Exit Configure VPLS mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 3.3.3.3/32 secondary | Configure IP address for the loopback interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface cd6/1 | Enter interface mode |
| (config-if)#ip address 30.1.1.1/24 | Configure IP address for the interface |
| (config-if)#label-switching | Configure label switching |
| (config-if)#exit | Exit interface mode |
| (config)#vpls fib-entry 100 peer 1.1.1.1 4000 cd6/1 3000 | Configure VPLS FIB entry for VPLS peer PE-1. |
| (config)#vpls fib-entry 100 peer 2.2.2.2 4500 cd6/1 3500 | Configure VPLS FIB entry for VPLS peer PE-2. |
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)# router-id 3.3.3.3 | Configure router id same as loopback IP address |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure transport address as loopback address. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |

| | |
|---|---|
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config)#router ospf 100 | Configure the routing process and specify the routing process ID(100). |
| (config-router)#network 30.30.1.0/24 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config-router)#network 3.3.3.3/32 area 0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the interface address. |
| (config)#commit | Commit the transaction. |

Validation

Enter the commands listed in the sections below to confirm the configurations.

Verify VPLS Session on PE-1

```
PE1#show mpls vpls detail
Virtual Private LAN Service Instance: static_vpls, ID: 100
SIG-Protocol: STATIC
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe2.1000
Status: Up
Subinterface Match Criteria(s) :
dot1q 100
Split-horizon group : access1

Mesh Peers:
2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 01:14:56)
Tunnel-Id: 11
3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 01:14:56)
Tunnel-Id: 22
```

```
PE1#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|---|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | | |
| 100 | 2.2.2.2 | 120 | 1001 | N/A | 2001 | 2/Up | 5 | S |
| TATIC | Active | 01:15:57 | - | | | | | |
| 100 | 3.3.3.3 | 103 | 3000 | N/A | 4000 | 2/Up | 6 | S |
| TATIC | Active | 01:15:57 | - | | | | | |

Verify VPLS Session on PE-2

```
PE2#show mpls vpls detail
```

```

Virtual Private LAN Service Instance: static_vpls, ID: 100
SIG-Protocol: STATIC
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
MAC Withdrawal:

```

```

Configured interfaces:
Interface: xe0.1000
Status: Up
Subinterface Match Criteria(s) :
dot1q 100
Split-horizon group : access1

```

```

Mesh Peers:
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 18:41:17)
Tunnel-Id: 11
3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:28:53)
Tunnel-Id: 33

```

```

PE2# show mpls vpls mesh

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|---------|-----------|--------------|----------|--------------|-----------|---------|----------|
| 100 | 1.1.1.1 | 201 | 2500 | N/A | 1500 | 2/Up | 3 S |
| TATIC | Active | 18:41:45 | - | | | | |
| 100 | 3.3.3.3 | 301 | 3500 | N/A | 4500 | 2/Up | 14 S |
| TATIC | Active | 00:29:21 | - | | | | |

Verify VPLS Session on PE-3

```

PE3# show mpls vpls mesh

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|---------|-----------|--------------|----------|--------------|-----------|---------|----------|
| 100 | 1.1.1.1 | 402 | 4001 | N/A | 3000 | 2/Up | 1 S |
| TATIC | Active | 00:35:23 | - | | | | |
| 100 | 2.2.2.2 | N/A | 4501 | N/A | 3500 | 0/Dn | 3 S |
| TATIC | Active | 00:36:45 | - | | | | |

```

PE3#sh mpls vpls detail
Virtual Private LAN Service Instance: static_vpls, ID: 100
SIG-Protocol: STATIC
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
MAC Withdrawal:

```

```

Configured interfaces:
Interface: cd4/4.1000
Status: Up
Subinterface Match Criteria(s) :
dot1q 100
Split-horizon group : access1

```

```
Mesh Peers:
 1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:36:05)
 Tunnel-Id: 11
 2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Dn) (Uptime: 00:45:16)
 Tunnel-Id: 2
```

Remove Configurations

Follow these steps to remove VPLS peer and VPLS spoke FIB entries from router PE-2.

| | |
|---|-------------------------------------|
| #configure terminal | Enter configure mode |
| (config)#no vpls fib-entry 100 peer 1.1.1.1 | Remove VPLS FIB for VPLS peer PE-1. |
| (config)#no vpls fib-entry 100 peer 3.3.3.3 | Remove VPLS FIB for VPLS peer PE-3. |
| (config)#exit | Exit Configure mode |

BGP-VPLS Configuration

This chapter contains configurations for VPLS with Border Gateway Protocol (BGP) Signaling.

Virtual Private LAN Service (VPLS) provides Ethernet-based multipoint-to-multipoint communication over IP/MPLS networks. It allows sites that are geographically separated to share an Ethernet broadcast domain by connecting sites through pseudowires. For instance, a set of Kompella circuits is grouped by a common VPLS identifier to achieve this service objective.

A Pseudowire (PW) consists of a pair of point-to-point, single-hop, unidirectional LSPs, placed in opposite directions, each identified by a PW label, also referred to as a Virtual Connection (VC) label.

BGP is used to signal the VCs and for auto-discovery of neighbors. A service provider may use either LDP or RSVP-TE or add static provisioning to set up LSP tunnels for data transport through virtual circuits.

The VPLS identifier is exchanged with the labels, and both PWs are linked and associated with a particular VPLS instance.



Note: In Inter-AS, OcNOS accepts information from any other AS but the same VPN-ID/VPLS-ID (*: VPLS-ID). OcNOS does not have explicit RD/RT (import/export) support for BGP VPLS. RD/RT are automatically generated based on the configured BGP AS number and VPN-ID/VPLS-ID as (AS-number: VPN-ID).




Note: In OcNOS, only the RD/RT format of AS:NN (Autonomous System:Number) is supported. The IP:NN (IP address:Number) format is not supported.



Note: If a transit Route Reflector (RR) node is reloaded, any BGP VPLS sessions traversing that RR will flap. As a result, the VPLS mesh VC FIB entry is removed and then reinstalled using the next available matching path.

Topology

The diagram depicts the topology for the configuration examples that follow.

Figure 75. Sample Topology for VPLS with BGP Signaling

BGP-VPLS Configuration

PE-1

| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface xe2 | Specify the interface (xe1) to be configured. |
| (config-if)#ip address 10.10.10.1/24 | Set the IP address of the interface to 10.10.10.1/24. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback address. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32. |
| (config-if)#exit | Exit interface mode |
| (config)#mpls vpls v1 25 | Create an instance of VPLS, and switch to the VPLS command mode, by specifying the VPLS name (v1) and VPLS ID (25). |
| (config-vpls)#vpls-mtu 1400 | Configure the MTU for the VPLS. (Default is 1500; range is <576 - 65535>. |
| (config-vpls)#signaling bgp | Enter the Signaling bgp mode for BGP VPLS. |
| (config-vpls-sig)#ve-id 1 | Configure VE ID, which is mandatory for BGP VPLS, otherwise, Signaling does not take place. VE ID should be unique per VPLS instance. |
| (config-vpls-sig)#exit | Exit is a mandatory command for signaling BGP configuration to take affect. If exit is not given BGP signaling does not take place. |
| (config-vpls)#exit | Exit VPLS mode. |
| (config)#interface xe1.25 switchport | Configure static sub-interface. |

| | |
|---|--|
| (config-if)#encapsulation dot1q 25 | Configure encapsulation. |
| (config-if)#split-horizon group access1 | Configure split-horizon group on VPLS. |
| (config-if)#access-if-vpls | Access VPLS under sub interface. |
| (config-acc-if-vpls)# mpls-vpls v1 | Bind the VPLS to the Access Interface. |
| (config-acc-if-vpls)#commit | Commit the transaction. |

PE1 - LDP

| | |
|--|---|
| #configure terminal | Enter configuration mode |
| (config)#router ldp | Enter Router LDP mode. |
| (config-router)#router-id 1.1.1.1 | Configure the router ID. |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure the transport address for a label space by binding the address to a loopback address. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#enable-ldp ipv4 | Enable LDP on interface xe2. |
| (config-if)#commit | Commit the transaction. |

PE1 - OSPF

| | |
|--|--|
| #configure terminal | Enter configure mode |
| (config)#router ospf 1 | Configure the OSPF routing process, and specify the process ID. |
| (config-router)#network 10.10.10.0/24 area 0 | Define the interfaces on which OSPF runs, and specify the backbone area 0. |
| (config-router)#network 1.1.1.1/32 area 0 | |
| (config-router)#commit | Commit the transaction. |

PE1 - BGP

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#router bgp 100 | Enter BGP Configure mode. |
| (config-router)#neighbor 2.2.2.2 remote-as 100 | Configure PE2 as an iBGP peer. |
| (config-router)#neighbor 2.2.2.2 update-source lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| (config-router)#neighbor 3.3.3.3 remote-as 100 | Configure PE3 as an iBGP peer. |
| (config-router)#neighbor 3.3.3.3 update-source lo | Update the source as loopback for iBGP peering with the remote PE3 router |
| (config-router)#address-family l2vpn vpls | Configure address-family L2VPN VPLS. |
| (config-router-af)#neighbor 2.2.2.2 activate | Activate PE2 in the VPLS address family. |
| (config-router-af)#neighbor 3.3.3.3 activate | Activate PE3 in the VPLS address family. |
| (config-router-af)#commit | Commit the transaction. |

PE2

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface xe7 | Specify the interface (xe7 to be configured. |
| (config-if)#ip address 20.20.20.1/24 | Set the IP address of the interface to 20.20.20.1/24. |
| (config-if)#label-switching | Enable label switching on interface xe7. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback address. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#mpls vpls v1 25 | Create an instance of VPLS, and switch to the VPLS command mode, by specifying the VPLS name (v1) and VPLS ID (25). |
| (config-vpls)#vpls-mtu 1400 | Configure the MTU for the VPLS. (Default is 1500; range is <576 - 65535>.) |
| (config-vpls)#signaling bgp | Enter the Signaling BGP mode for BGP VPLS. |
| (config-vpls-sig)#ve-id 2 | Configure ve-id, which is mandatory for BGP VPLS. Without a ve-id Signaling does not take place. VE ID should be unique per VPLS instance. |
| (config-vpls-sig)#exit | Exit is a mandatory command for signaling BGP configuration to take affect. If exit is done, BGP signaling does not take place. |
| (config-vpls)#exit | Exit VPLS mode. |
| (config)#interface xe26.25 | Specify the interface (xe26.25) to be configured. |
| (config-if)#switchport | Switch to Layer-2 mode. (VPLS can only be bound on the Layer-2 port.) |
| (config-if)#encapsulation dot1q 25 | Configure the encapsulation. |
| (config-if-vpls)#split-horizon group access1 | Configure split-horizon group on VPLS |
| (config-if)#access-if-vpls | Access VPLS under sub interface |
| (config-acc-if-vpls)# mpls-vpls v1 | Bind the VPLS to the Access Interface. |
| (config-if)#commit | Commit the transaction. |

PE2 - LDP

| | |
|--|---|
| #configure terminal | Enter configuration mode |
| (config)#router ldp | Enter Router LDP mode. |
| (config-router)#router-id 2.2.2.2 | Configure the router ID. |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address for a label space by binding the address to a loopback address. |

| | |
|-----------------------------|---|
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe7 | Specify the interface (xe7) to be configured. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the specified interface (xe7). |
| (config-if)#commit | Commit the transaction. |

PE2 - OSPF

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#router ospf 1 | Configure the OSPF routing process, and specify the process ID. |
| (config-router)#network 20.20.20.0/24 area 0 | Define the interfaces on which OSPF runs, and specify the backbone area 0. |
| (config-router)#network 2.2.2.2/32 area 0 | |
| (config-router)#commit | Commit the transaction. |

PE2 - BGP

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#router bgp 100 | Enter BGP router mode. |
| (config-router)#neighbor 1.1.1.1 remote-as 100 | Configure PE1 as an iBGP peer. |
| (config-router)#neighbor 1.1.1.1 update-source lo | Update the source as loopback for iBGP peering with the remote PE1 router. |
| (config-router)#neighbor 3.3.3.3 remote-as 100 | Configure PE3 as an iBGP peer. |
| (config-router)#neighbor 3.3.3.3 update-source lo | Update the source as loopback for iBGP peering with the remote PE3 router. |
| (config-router)#address-family l2vpn vpls | Configure address-family L2VPN VPLS. |
| (config-router-af)#neighbor 1.1.1.1 activate | Activate PE1 in the VPLS address family. |
| (config-router-af)#neighbor 3.3.3.3 activate | Activate PE3 in the VPLS address family. |
| (config-router-af)#commit | Commit the transaction. |

PE3

| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface xe4 | Specify the interface (xe4) to be configured. |
| (config-if)#ip address 30.30.30.1/24 | Set the IP address of the interface to 30.30.30.1/24. |
| (config-if)#label-switching | Enable label switching on interface xe4. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface lo | Specify the loopback address. |
| (config-if)#ip address 3.3.3.3/32 secondary | Set the IP address of the loopback interface to 3.3.3.3/32. |
| (config-if)#exit | Exit interface mode. |

| | |
|---|---|
| <code>(config)#mpls vpls v1 25</code> | Create an instance of VPLS, and switch to the VPLS command mode, by indicating the VPLS name (v1) and VPLS ID (25). |
| <code>(config-vpls)#vpls-mtu 1400</code> | Configure the MTU for the VPLS. Default is 1500; range is <576 - 65535>. |
| <code>(config-vpls)#signaling bgp</code> | Enter the Signaling BGP mode, for BGP VPLS. |
| <code>(config-vpls-sig)#ve-id 3</code> | Configure ve-id, which is mandatory for BGP VPLS. Without a ve-id Signaling does not take place. VE ID should be unique per VPLS instance |
| <code>(config-vpls-sig)#exit</code> | Exit is a mandatory command for signaling BGP configuration to take affect. If exit is not done, BGP signaling does not take place. |
| <code>(config-vpls)#exit</code> | Exit VPLS mode. |
| <code>(config)#interface xe25.25</code> | Specify the interface (xe25.25) to be configured. |
| <code>(config-if)#switchport</code> | Switch to Layer-2 mode. (VPLS can be bound only on the Layer-2 port.) |
| <code>(config-if)#encapsulation dot1q 25</code> | Configure encapsulation under a subinterface |
| <code>(config-if-vpls)#split-horizon group access1</code> | Configure split-horizon group on VPLS |
| <code>(config-if)#access-if-vpls</code> | Access VPLS under sub interface. |
| <code>(config-acc-if-vpls)# mpls-vpls v1</code> | Bind the VPLS to the Access Interface. |
| <code>(config-if)#commit</code> | Commit the transaction. |



Note: VE ID's range is from 1 to 64. Administrator should configure the VE ID's accordingly in their Network.

PE3 - LDP

| | |
|---|---|
| <code>#configure terminal</code> | Enter configuration mode. |
| <code>(config)#router ldp</code> | Enter Router LDP mode. |
| <code>(config-router)#router-id 3.3.3.3</code> | Configure the router ID. |
| <code>(config-router)#transport-address ipv4 3.3.3.3</code> | Configure the transport address for a label space by binding the address to a loopback address. |
| <code>(config-router)#exit</code> | Exit Router mode. |
| <code>(config)#interface xe4</code> | Specify the interface (xe4) to be configured. |
| <code>(config-if)#enable-ldp ipv4</code> | Enable LDP on the interface. |
| <code>(config-if)#commit</code> | Commit the transaction. |

PE3 - OSPF

| | |
|-------------------------------------|---|
| <code>#configure terminal</code> | Enter configuration mode. |
| <code>(config)#router ospf 1</code> | Configure the OSPF routing process, and specify |

| | |
|--|--|
| | the process ID. |
| (config-router)#network 30.30.30.0/24 area 0 | Define the interfaces on which OSPF runs, and specify the backbone area 0. |
| (config-router)#network 3.3.3.3/32 area 0 | |
| (config-router)#commit | Commit the transaction. |

PE3 - BGP

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#router bgp 100 | Enter BGP Router mode. |
| (config-router)#neighbor 1.1.1.1 remote-as 100 | Configure PE1 as an iBGP peer. |
| (config-router)#neighbor 1.1.1.1 update-source lo | Update the source as loopback for iBGP peering with the remote PE1 router. |
| (config-router)#neighbor 2.2.2.2 remote-as 100 | Configure PE2 as an iBGP peer. |
| (config-router)#neighbor 2.2.2.2 update-source lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| (config-router)#address-family l2vpn vpls | Configure address-family L2VPN VPLS. |
| (config-router-af)#neighbor 1.1.1.1 activate | Activate PE1 in the VPLS address family. |
| (config-router-af)#neighbor 2.2.2.2 activate | Activate PE2 in the VPLS address family. |
| (config-router-af)#commit | Commit the transaction. |

P

| | |
|--------------------------------------|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#ip address 10.10.10.2/24 | Set the IP address of the interface to 10.10.10.2/24. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe4 | Specify the interface (xe4) to be configured. |
| (config-if)#ip address 30.30.30.2/24 | Set the IP address of the interface to 30.30.30.2/24. |
| (config-if)#label-switching | Enable label switching on interface xe4. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe4 | Specify the interface (xe4) to be configured. |
| (config-if)#ip address 20.20.20.2/24 | Set the IP address of the loopback interface to 20.20.20.2/24. |
| (config-if)#label-switching | Enable label switching on interface xe4. |
| (config-if)#commit | Commit the transaction. |

P - LDP

| | |
|-----------------------------|---|
| #configure terminal | Enter configuration mode. |
| (config)#router ldp | Enter Router LDP mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#interface xe2 | Specify the interface (xe2) to be configured. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe4 | Specify the interface (xe4) to be configured. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe7 | Specify the interface (xe7) to be configured. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the interface. |
| (config-if)#commit | Commit the transaction. |

P - OSPF

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#router ospf 1 | Configure the OSPF routing process, and specify the process ID. |
| (config-router)#network 10.10.10.0/24 area 0 | Define the interfaces on which OSPF runs, and specify the backbone area 0. |
| (config-router)#network 20.20.20.0/24 area 0 | |
| (config-router)#network 30.30.30.0/24 area 0 | |
| (config-router)#commit | Commit the transaction. |

Validation**P1**

```
P1#show ip ospf neighbor
Total number of full neighbors: 3
OSPF process 100 VRF(default):
Neighbor ID    Pri   State       Dead Time   Address        Interface      Instance ID
1.1.1.1        1     Full/Backup  00:00:30    10.10.10.1     xe2            0
2.2.2.2        1     Full/Backup  00:00:37    20.20.20.1     xe7            0
3.3.3.3        1     Full/Backup  00:00:30    30.30.30.1     xe4            0

P1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually

Code Peer IP Address    IF Name    My Role    State        KeepAlive  UpTime
-----
2.2.2.2                xe7        Active     OPERATIONAL 30         22:14:11
1.1.1.1                xe2        Active     OPERATIONAL 30         20:11:39
3.3.3.3                xe4        Active     OPERATIONAL 30         20:05:31
```


PE1

```
PE1#sh run vpls
```

```
mpls vpls vl 25
  vpls-mtu 1400
  signaling bgp
  ve-id 1
  exit-signaling
  exit-vpls
interface xel.25 switchport
  split-horizon group access1
  access-if-vpls
  mpls-vpls vl
```

```
PE1#sh ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 4.4.4.4 | xe2 | Passive | OPERATIONAL | 30 | 19:16:49 |

```
PE1#show mpls vpls detail
```

```
Virtual Private LAN Service Instance: vl, ID: 25
SIG-Protocol: BGP
  Route-Distinguisher :100:25
  Route-Target :100:25
  VE-ID :1
  Attachment-Circuit: UP
  Learning: Enabled
  Control-Word: Disabled
  Flow Label Status: Disabled, Direction: None, Static: No
  Group ID: 0, Configured MTU: 1400
  Description: none
  service-tpid: dot1q
  Operating mode: Raw
  MAC Withdrawal:
```

```
Configured interfaces:
```

```
  Interface: xel.25
  Status: Up
  Subinterface Match Criteria(s) :
  dot1q 25
  Split-horizon group : access1
```

```
Mesh Peers:
```

```
  2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 01:25:35)
  3.3.3.3 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 01:06:50)
```

```
PE1#sh mpls vpls mesh
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 25 | 2.2.2.2 | 24323 | 26242 | xe2 | 25600 | 2/Up | 1 B |
| GP | Active | 01:26:34 | - | | | | |
| 25 | 3.3.3.3 | 24325 | 26241 | xe2 | 25600 | 2/Up | 2 B |
| GP | Active | 01:07:49 | - | | | | |

```
PE1#sh bgp l2vpn vpls detail
```

```
VPLS ID: 25
VE-ID: 1
Discovered Peers: 2
```

```

Route-Target: 100:25
Local RD: 100:25
All Local Label Blocks:
  [LB:26240, VBO:1, VBS:64]
Mesh Peers:
  BGP Peer:2.2.2.2/32
    VC Nbr Address:2.2.2.2, RD:100:25, VE-ID:3
    VC Details: VC-ID:13
    Local MTU:1400, Remote MTU:1400
    Remote (LB:25600,VBO:1,VBS:64)  Local (LB:26240,VBO:1,VBS:64)
    LB sent on known VEID:Yes
    In Label:26242, Out Label:25600
    PW Status:Established
    VC Installed:Yes
    VC Signaled Time: 01:27:54
    Extended-Community Color:0

  BGP Peer:3.3.3.3/32
    VC Nbr Address:3.3.3.3, RD:100:25, VE-ID:2
    VC Details: VC-ID:12
    Local MTU:1400, Remote MTU:1400
    Remote (LB:25600,VBO:1,VBS:64)  Local (LB:26240,VBO:1,VBS:64)
    LB sent on known VEID:Yes
    In Label:26241, Out Label:25600
    PW Status:Established
    VC Installed:Yes
    VC Signaled Time: 01:09:09
    Extended-Community Color:0

VPLS ID: 25
VE-ID: 1
Discovered Peers: 2
Route-Target: 100:25
Local RD: 100:25
All Local Label Blocks:
  [LB:26240, VBO:1, VBS:64]
Mesh Peers:
  BGP Peer:2.2.2.2/32
    VC Nbr Address:2.2.2.2, RD:100:25, VE-ID:3
    VC Details: VC-ID:13
    Local MTU:1400, Remote MTU:1400
    Remote (LB:25600,VBO:1,VBS:64)  Local (LB:26240,VBO:1,VBS:64)
    LB sent on known VEID:Yes
    In Label:26242, Out Label:25600
    PW Status:Established
    VC Installed:Yes
    VC Signaled Time: 01:29:06
    Extended-Community Color:0

  BGP Peer:3.3.3.3/32
    VC Nbr Address:3.3.3.3, RD:100:25, VE-ID:2
    VC Details: VC-ID:12
    Local MTU:1400, Remote MTU:1400
    Remote (LB:25600,VBO:1,VBS:64)  Local (LB:26240,VBO:1,VBS:64)
    LB sent on known VEID:Yes
    In Label:26241, Out Label:25600
    PW Status:Established
    VC Installed:Yes
    VC Signaled Time: 01:10:21
    Extended-Community Color:0

PE1#sh bgp l2vpn vpls summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|----|--------|--------|--------|-----|------|---------|--------------|------|
|----------|---|----|--------|--------|--------|-----|------|---------|--------------|------|

| | | | | | | | | | |
|---------|---|-----|-----|-----|---|---|---|----------|---|
| 2.2.2.2 | 4 | 100 | 214 | 212 | 1 | 0 | 0 | 01:29:39 | 1 |
| 3.3.3.3 | 4 | 100 | 336 | 534 | 1 | 0 | 0 | 01:10:55 | 1 |

Total number of neighbors 2

Total number of Established sessions 2

PE2

```
PE2#show ldp session
```

```
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 4.4.4.4 | xe7 | Passive | OPERATIONAL | 30 | 21:14:46 |

```
PE2#show run vpls
```

```
mpls vpls vl 25
  vpls-mtu 1400
  signaling bgp
  ve-id 3
  exit-signaling
  exit-vpls
interface xe26.25 switchport
  split-horizon group access1
  access-if-vpls
  mpls-vpls vl
```

```
PE2#show mpls vpls mesh
```

```
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 25 | 1.1.1.1 | 24322 | 25600 | xe7 | 26242 | 2/Up | 1 B |
| GP | Active | 01:34:16 | - | | | | |
| 25 | 3.3.3.3 | 24324 | 25601 | xe7 | 25602 | 2/Up | 2 B |
| GP | Active | 01:15:34 | - | | | | |

```
PE2#sh bgp l2vpn vpls
```

| VPLS-ID | VE-ID | Discovered-Peers | Route-Target |
|---------|-------|------------------|--------------|
| 25 | 3 | 2 | 100:25 |

```
PE2#show bgp l2vpn vpls detail
```

```
VPLS ID: 25
```

```
VE-ID: 3
```

```
Discovered Peers: 2
```

```
Route-Target: 100:25
```

```
Local RD: 100:25
```

```
All Local Label Blocks:
```

```
[LB:25600, VBO:1, VBS:64]
```

```
Mesh Peers:
```

```
BGP Peer:1.1.1.1/32
```

```
VC Nbr Address:1.1.1.1, RD:100:25, VE-ID:1
```

```
VC Details: VC-ID:31
```

```
Local MTU:1400, Remote MTU:1400
```

```
Remote (LB:26240,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
```

```
LB sent on known VEID:Yes
```

```
In Label:25600, Out Label:26242
```

```
PW Status:Established
```

```
VC Installed:Yes
```

```
VC Signaled Time: 01:34:16
```

```
Extended-Community Color:0
```

```
BGP Peer:3.3.3.3/32
```

```
VC Nbr Address:3.3.3.3, RD:100:25, VE-ID:2
```

```
VC Details: VC-ID:32
```

```

Local MTU:1400, Remote MTU:1400
Remote (LB:25600,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25601, Out Label:25602
PW Status:Established
VC Installed:Yes
VC Signaled Time: 01:15:34
Extended-Community Color:0

PE2#show bgp l2vpn vpls 25
VPLS ID: 25
VE-ID: 3
Discovered Peers: 2
Route-Target: 100:25
Local RD: 100:25
All Local Label Blocks:
[LB:25600, VBO:1, VBS:64]
Mesh Peers:
BGP Peer:1.1.1.1/32
VC Nbr Address:1.1.1.1, RD:100:25, VE-ID:1
VC Details: VC-ID:31
Local MTU:1400, Remote MTU:1400
Remote (LB:26240,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25600, Out Label:26242
PW Status:Established
VC Installed:Yes
VC Signaled Time: 01:34:16
Extended-Community Color:0

BGP Peer:3.3.3.3/32
VC Nbr Address:3.3.3.3, RD:100:25, VE-ID:2
VC Details: VC-ID:32
Local MTU:1400, Remote MTU:1400
Remote (LB:25600,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25601, Out Label:25602
PW Status:Established
VC Installed:Yes
VC Signaled Time: 01:15:34
Extended-Community Color:0

PE2#show bgp l2vpn vpls summary
BGP router identifier 2.2.2.2, local AS number 100
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries
Neighbor      V    AS      MsgRcv   MsgSen   TblVer    InQ    OutQ    Up/Down   State/PfxRcd   Desc
1.1.1.1        4    100       2717     2720      1         0       0    01:34:23         1
3.3.3.3        4    100       2722     2726      1         0       0    01:15:42         1
Total number of neighbors 2
Total number of Established sessions 2

```

PE3

```

PE3#show run vpls

mpls vpls v1 25
vpls-mtu 1400
signaling bgp
ve-id 2
exit-signaling
exit-vpls

interface xe25.25 switchport
split-horizon group access1
access-if-vpls
mpls-vpls v1

```

```
PE3#show mpls vpls mesh
```

```
(m) - Service mapped over multipath transport
```

```
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 25 | 1.1.1.1 | 24326 | 25600 | xe4 | 26241 | 2/Up | 4 B |
| GP | Active | 01:13:19 | - | | | | |
| 25 | 2.2.2.2 | 24327 | 25602 | xe4 | 25601 | 2/Up | 3 B |
| GP | Active | 01:13:22 | - | | | | |

```
PE3#sh bgp l2vpn vpls
```

| VPLS-ID | VE-ID | Discovered-Peers | Route-Target |
|---------|-------|------------------|--------------|
| 25 | 2 | 2 | 100:25 |

```
PE3#show bgp l2vpn vpls detail
```

```
VPLS ID: 25
```

```
VE-ID: 2
```

```
Discovered Peers: 2
```

```
Route-Target: 100:25
```

```
Local RD: 100:25
```

```
All Local Label Blocks:
```

```
[LB:25600, VBO:1, VBS:64]
```

```
Mesh Peers:
```

```
BGP Peer:1.1.1.1/32
```

```
VC Nbr Address:1.1.1.1, RD:100:25, VE-ID:1
```

```
VC Details: VC-ID:21
```

```
Local MTU:1400, Remote MTU:1400
```

```
Remote (LB:26240,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
```

```
LB sent on known VEID:Yes
```

```
In Label:25600, Out Label:26241
```

```
PW Status:Established
```

```
VC Installed:Yes
```

```
VC Signaled Time: 01:13:19
```

```
Extended-Community Color:0
```

```
BGP Peer:2.2.2.2/32
```

```
VC Nbr Address:2.2.2.2, RD:100:25, VE-ID:3
```

```
VC Details: VC-ID:23
```

```
Local MTU:1400, Remote MTU:1400
```

```
Remote (LB:25600,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
```

```
LB sent on known VEID:Yes
```

```
In Label:25602, Out Label:25601
```

```
PW Status:Established
```

```
VC Installed:Yes
```

```
VC Signaled Time: 01:13:22
```

```
Extended-Community Color:0
```

```
P3#show bgp l2vpn vpls 25
```

```
VPLS ID: 25
```

```
VE-ID: 2
```

```
Discovered Peers: 2
```

```
Route-Target: 100:25
```

```
Local RD: 100:25
```

```
All Local Label Blocks:
```

```
[LB:25600, VBO:1, VBS:64]
```

```
Mesh Peers:
```

```
BGP Peer:1.1.1.1/32
```

```
VC Nbr Address:1.1.1.1, RD:100:25, VE-ID:1
```

```
VC Details: VC-ID:21
```

```
Local MTU:1400, Remote MTU:1400
```

```
Remote (LB:26240,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
```

```
LB sent on known VEID:Yes
```

```

In Label:25600, Out Label:26241
PW Status:Established
VC Installed:Yes
VC Signaled Time: 01:13:19
Extended-Community Color:0

BGP Peer:2.2.2.2/32
VC Nbr Address:2.2.2.2, RD:100:25, VE-ID:3
VC Details: VC-ID:23
Local MTU:1400, Remote MTU:1400
Remote (LB:25600,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25602, Out Label:25601
PW Status:Established
VC Installed:Yes
VC Signaled Time: 01:13:22
Extended-Community Color:0

PE3#show bgp l2vpn vpls summary
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS    MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd   Desc
1.1.1.1            4   100      217     256      1      0      0   01:13:55         1
2.2.2.2            4   100      182     185      1      0      0   01:13:58         1

Total number of neighbors 2

Total number of Established sessions 2

PE3#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
  4.4.4.4                xe4        Passive    OPERATIONAL  30      19:08:09

```

BGP Four Bytes ASN with VPLS Configuration

Overview

In case of BGP-VPLS with four bytes ASN, RD will have administrator field expanded to accommodate 4 byte ASN whereas assigned number field remains two bytes. This adjustment allows the RD to continue providing unique identification for VPLS instances.



Note:

Enabling or disabling the four bytes ASN capability (bgp extended-asn-cap) requires the removal of certain existing configurations, such as VRFs associated with L3VPN and EVPN, as well as BGP auto-discovery configuration used in LDP-VPLS.

Prerequisites

Configure Loopback Interfaces and Core Interfaces with IP Addresses

Ensure that the loopback and core interfaces are configured with IPv4 addresses.

```
interface lo
  ip address 127.0.0.1/8
  ip address 30.0.1.2/32 secondary
  ipv6 address ::1/128

interface xe4
  ip address 10.0.1.12/31
```

Configure IGP for Dynamic Routing

OSPF Configuration

Configure OSPF with router ID and advertise the networks. Set the interface as an OSPF point-to-point network type.



Note: Either ISIS or OSPF can be used as an IGP protocol.

```
router ospf 100
  ospf router-id 30.0.1.2
  network 10.0.1.12/31 area 0.0.0.0
  network 30.0.1.2/32 area 0.0.0.0

interface xe4
  ip ospf network point-to-point
```

Configure Transport Protocol

Enable RSVP globally using the router rsvp command, and configure the interface with label-switching and RSVP support using label-switching and enable-rsvp commands.



Note: Either RSVP or LDP can be used for transport.

```
router rsvp

Interface xe4
  label-switching
  enable-rsvp
```

Topology

The topology shows an MPLS Layer 2 VPN setup where CE1 and CE2 connect via PE1 and PE2, both in AS 80000. Core router P1 performs label switching, with /31 point-to-point links and /32 loopbacks used for routing and MPLS label distribution.




Figure 76. BGP 4-bytes ASN with VPLS

Configuration

This configuration sets up MPLS Layer 2 VPN using RSVP and OSPF for IGP routing. BGP is configured for VPLS signaling and interfaces are prepared for MPLS label switching and customer VLAN access. The setup enables end-to-end connectivity between CE1 and CE2 over VPLS service using BGP 4 byte ASN.

Configure PE1 router as follows

1. Enable resource reservation protocol (RSVP) in global configuration mode.

```
PE1(config)#router rsvp
```

2. Configure a IP address on the loopback interface. This address will be used as the router ID and for BGP peering.

```
PE1(config)#interface lo
PE1(config-if)#ip address 30.0.1.1/32 secondary
PE1(config-if)#commit
```

3. Assign an IP address, enable MPLS label switching, define the OSPF network type, and enable RSVP on the interface.

```
PE1(config)#interface xe29
PE1(config-if)#ip address 10.0.1.20/31
PE1(config-if)#label-switching
PE1(config-if)#ip ospf network point-to-point
PE1(config-if)#enable-rsvp
PE1(config-if)#commit
```

4. Configure OSPF with a specific process ID, set the router ID to the loopback address, and advertise both the core interface address and the loopback address into area 0.

```
PE1(config)#router ospf 100
PE1(config-router)#ospf router-id 30.0.1.1
```



```
PE1(config-router)#network 10.0.1.20/31 area 0.0.0.0
PE1(config-router)#network 30.0.1.1/32 area 0.0.0.0
PE1(config-if)#commit
```

5. Create an RSVP trunk that defines the remote endpoint (typically the loopback address of PE2). This trunk allows RSVP signaling to establish LSPs between PEs.

```
PE1(config)#rsvp-trunk PE1-PE2
PE1(config-trunk)#to 30.0.1.2
PE1(config-if)#commit
```

6. Enable the 4-byte ASN capability and configure BGP. Define the neighbor using the remote PE's loopback address, specify the ASN, and enable both IPv4 unicast and L2VPN VPLS address families.

```
PE1(config)#bgp extended-asn-cap
PE1(config)#router bgp 80000
PE1(config-router)#neighbor 30.0.1.2 remote-as 80000
PE1(config-router)#neighbor 30.0.1.2 update-source lo
PE1(config-router)#commit
PE1(config-router)#address-family ipv4 unicast
PE1(config-router-af)#neighbor 30.0.1.2 activate
PE1(config-router-af)#commit
PE1(config-router-af)#exit
PE1(config-router)#address-family l2vpn vpls
PE1(config-router-af)#neighbor 30.0.1.2 activate
PE1(config-router-af)#commit
```

7. Create a VPLS instance and configure BGP as the signaling protocol. The ve-id uniquely identifies this PE in the VPLS instance.

```
PE1(config)#mpls vpls PE1-PE2 101
PE1(config-vpls)#signaling bgp
PE1(config-vpls-sig)#ve-id 10
PE1(config-vpls-sig)#commit
PE1(config-vpls-sig)#exit
```

8. Set up the access (customer-facing) interface with VLAN encapsulation and bind it to the VPLS instance.

```
PE1(config)#interface xe27.101 switchport
PE1(config-if)#encapsulation dot1q 101
PE1(config-if)#access-if-vpls
PE1(config-acc-if-vpls)#mpls-vpls PE1-PE2
PE1(config-acc-if-vpls)#commit
PE1(config-acc-if-vpls)#exit
```

Configure P1 router as follows

1. Enable resource reservation protocol (RSVP) in global configuration mode.

```
P1(config)# router rsvp
```

2. Enter the loopback interface configuration mode and assign a IP address.

```
P1(config)# interface lo
P1(config-if)# ip address 30.0.1.3/32 secondary
```

3. Commit and exit the interface configuration

```
P1(config-if)# commit
P1(config-if)# exit
```

4. Assign an IP address, enable MPLS label switching, define the OSPF network type, and enable RSVP on the core interfaces.

```
P1(config)# interface xe2
P1(config-if)# ip address 10.0.1.21/31
P1(config-if)# label-switching
```

```
P1(config-if)# ip ospf network point-to-point
P1(config-if)# enable-rsvp
P1(config-if)# commit
P1(config-if)# exit
P1(config)# interface xe4
P1(config-if)# ip address 10.0.1.13/31
P1(config-if)# label-switching
P1(config-if)# ip ospf network point-to-point
P1(config-if)# enable-rsvp
P1(config-if)# commit
P1(config-if)# exit
```

5. Configure OSPF with a specific process ID, set the router ID to the loopback address, and advertise both the core interface address and the loopback address into area 0.

```
P1(config)# router ospf 100
P1(config-router)# ospf router-id 30.0.1.3
P1(config-router)# network 10.0.1.21/31 area 0.0.0.0
P1(config-router)# network 10.0.1.13/31 area 0.0.0.0
P1(config-router)# network 30.0.1.3/32 area 0.0.0.0
```

Configure PE2 router as follows

1. Enable resource reservation protocol (RSVP) in global configuration mode.

```
P1(config)# router rsvp
```

2. Configure Loopback Interface.

```
PE2(config)# interface lo
PE2(config-if)# ip address 30.0.1.2/32 secondary
PE2(config-if)# commit
PE2(config-if)# exit
```

3. Configure IP address on core interface xe4.

```
PE2(config)# interface xe4
PE2(config-if)# ip address 10.0.1.12/31
PE2(config-if)# label-switching
PE2(config-if)# ip ospf network point-to-point
PE2(config-if)# enable-rsvp
PE2(config-if)# commit
PE2(config-if)# exit
```

4. Configure OSPF with a specific process ID, set the router ID to the loopback address, and advertise both the core interface address and the loopback address into area 0.

```
PE2(config)# router ospf 100
PE2(config-router)# ospf router-id 30.0.1.2
PE2(config-router)# network 10.0.1.12/31 area 0.0.0.0
PE2(config-router)# network 30.0.1.2/32 area 0.0.0.0
PE2(config-router)# commit
PE2(config-router)# exit
```

5. Create an RSVP trunk that defines the remote endpoint (typically the loopback address of PE1). This trunk allows RSVP signaling to establish LSPs between PEs.

```
PE2(config)# rsvp-trunk PE2-PE1
PE2(config-trunk)# to 30.0.1.1
PE2(config-trunk)# commit
PE2(config-trunk)# exit
```

6. Enable the 4-byte ASN capability and configure BGP. Define the neighbor using the remote PE's loopback address, specify the ASN, and enable both IPv4 unicast and L2VPN VPLS address families.

```
PE2(config)# bgp extended-asn-cap
PE2(config)# router bgp 80000
PE2(config-router)# neighbor 30.0.1.1 remote-as 80000
```

```

PE2(config-router)# neighbor 30.0.1.1 update-source lo
PE2(config-router)# address-family ipv4 unicast
PE2(config-router-af)# neighbor 30.0.1.1 activate
PE2(config-router-af)# commit
PE2(config-router-af)# exit
PE2(config-router)# address-family l2vpn vpls
PE2(config-router-af)# neighbor 30.0.1.1 activate
PE2(config-router-af)# commit
PE2(config-router-af)# exit

```

7. Create a VPLS instance and configure BGP as the signaling protocol. The ve-id uniquely identifies this PE in the VPLS instance.

```

PE2(config-router)# mpls vpls PE2-PE1 101
PE2(config-vpls)# signaling bgp
PE2(config-vpls-sig)# ve-id 80
PE2(config-vpls-sig)# commit
PE2(config-vpls-sig)# exit

```

8. Set up the access (customer-facing) interface with VLAN encapsulation and bind it to the VPLS instance.

```

PE2(config)# interface xe10.101 switchport
PE2(config-if)# encapsulation dot1q 101
PE2(config-if)# access-if-vpls
PE2(config-acc-if-vpls)# mpls-vpls PE2-PE1
PE2(config-acc-if-vpls)# commit
PE2(config-acc-if-vpls)# exit

```

Validation

PE1



Note: OSPF adjacency should be in “Full” state.

```
PE1#show ip ospf neighbor
```

```
Total number of full neighbors: 1
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|-----------|-----------|-------------|
| 30.0.1.3 | 1 | Full/ - | 00:00:36 | 10.0.1.21 | xe29 | 0 |

```
PE1#
```

P1

```
P1#show ip ospf neighbor
```

```
Total number of full neighbors: 2
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|-----------|-----------|-------------|
| 30.0.1.2 | 1 | Full/ - | 00:00:38 | 10.0.1.12 | xe4 | 0 |
| 30.0.1.1 | 1 | Full/ - | 00:00:31 | 10.0.1.20 | xe2 | 0 |

PE2

```
PE2#show ip ospf neighbor
```

```
Total number of full neighbors: 1
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|-----------|-----------|-------------|
| 30.0.1.3 | 1 | Full/ - | 00:00:33 | 10.0.1.13 | xe4 | 0 |

```
PE2#
```

To verify the BGP session status between PE1 and PE2, ensure that the session is in the "Established" state. Additionally, in the output of the show ip bgp neighbors command, confirm that the 4-Octet ASN capability and the L2VPN VPLS capability are both advertised and received, indicating successful capability negotiation between the peers.

PE1

```

PE1#show ip bgp summary
BGP router identifier 30.0.1.1, local AS number 80000
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS      MsgRcv   MsgSen  TblVer    InQ   OutQ   Up/Down   State/PfxRcd   Desc
30.0.1.2          4  80000        5        5         1     0     0   00:00:30      0

Total number of neighbors 1

Total number of Established sessions 1
PE1#
PE1#show ip bgp neighbors
BGP neighbor is 30.0.1.2, remote AS 80000, local AS 80000, internal link, peer index: 2
  BGP version 4, local router ID 30.0.1.1, remote router ID 30.0.1.2
  BGP state = Established, up for 00:00:32
  Last read 00:00:06, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    4-Octet ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family L2VPN VPLS: advertised and received
  Received 5 messages, 0 notifications, 0 in queue
  Sent 6 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo

For address family: IPv4 Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  0 announced prefixes

For address family: L2VPN VPLS  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  1 announced prefixes

Connections established 1; dropped 0
Local host: 30.0.1.1, Local port: 179
Foreign host: 30.0.1.2, Foreign port: 54785
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460), Receive TCP MSS: (536)
Sock FD : (28)
Nexthop: 30.0.1.1 lo
Nexthop global: :: lo
Nexthop local: :: lo
BGP connection: non shared network

PE1#

```

PE2

```

PE2#show ip bgp summary
BGP router identifier 30.0.1.2, local AS number 80000
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

Neighbor      V    AS      MsgRcv   MsgSen  TblVer    InQ    OutQ    Up/Down  State/PfxRcd  Desc
30.0.1.1      4  80000      7        7        1       0       0    00:00:56      0

Total number of neighbors 1

Total number of Established sessions 1
PE2#
PE2#show ip bgp neighbors
BGP neighbor is 30.0.1.1, remote AS 80000, local AS 80000, internal link, peer index: 2
  BGP version 4, local router ID 30.0.1.2, remote router ID 30.0.1.1
  BGP state = Established, up for 00:01:00
  Last read 00:00:06, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    4-Octet ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Address family L2VPN VPLS: advertised and received
  Received 7 messages, 0 notifications, 0 in queue
  Sent 7 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo

For address family: IPv4 Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  0 announced prefixes

For address family: L2VPN VPLS  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  1 announced prefixes

Connections established 1; dropped 0
Local host: 30.0.1.2, Local port: 54785
Foreign host: 30.0.1.1, Foreign port: 179
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460), Receive TCP MSS: (536)
Sock FD : (28)
Nexthop: 30.0.1.2 lo
Nexthop global: :: lo
Nexthop local: :: lo
BGP connection: non shared network

PE2#

```

All RSVP sessions should be in “Up” state.

PE1

```

PE1#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

```

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
30.0.1.2    30.0.1.1    5001     2201     PRI     PE1-PE2-         24321
Primary                                -
Total 1 displayed, Up 1, Down 0.

Egress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
30.0.1.1    30.0.1.2    5001     2201     PRI     PE2-PE1-         24320
Primary                                -
Total 1 displayed, Up 1, Down 0.

PE1#

```

PE2

```

PE2#show rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
30.0.1.1    30.0.1.2    5001     2201     PRI     PE2-PE1-         24320
Primary                                -
Total 1 displayed, Up 1, Down 0.

Egress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime     Rt   Style  Labelin  Labelout
30.0.1.2    30.0.1.1    5001     2201     PRI     PE1-PE2-         24320
Primary                                -
Total 1 displayed, Up 1, Down 0.

PE2#

```

To verify BGP VPLS service status, the **show mpls vpls mesh** command should display the BGP VPLS session in the active state. The **show mpls vpls detail** output must show the Route Distinguisher (RD) and Route Targets (RTs) in the correct <4-byte ASN>:<VPLS ID> format. Additionally, the **show bgp l2vpn vpls detail** command should reflect accurate local and remote VE-IDs, proper RD information, and valid incoming and outgoing labels, ensuring proper signaling and label exchange between VPLS peers.

PE1

```

PE1#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID    Peer Addr      Tunnel-Label  In-Label    Network-Intf  Out-Label  Lkps/St  PW-INDEX
SIG-Protocol  Status      UpTime      Ext-Color
101          30.0.1.2    24321      25039      xe29          25033      2/Up      1          B
GP          Active      00:01:01    -
PE1#show mpls vpls detail
Virtual Private LAN Service Instance: PE1-PE2, ID: 101
SIG-Protocol: BGP
Route-Distinguisher :80000:101
Route-Target :80000:101
VE-ID :10
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled

```

```

Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe27.101
Status: Up
Subinterface Match Criteria(s) :
dot1q 101

Mesh Peers:
30.0.1.2 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:01:03)

```

```
PE1#show bgp l2vpn vpls detail
```

```

VPLS ID: 101
VE-ID: 10
Discovered Peers: 1
Route-Target: 80000:101
Local RD: 80000:101
All Local Label Blocks:
[LB:24960, VBO:1, VBS:64]
[LB:25024, VBO:65, VBS:64]
Mesh Peers:
BGP Peer:30.0.1.2/32
VC Nbr Address:30.0.1.2, RD:80000:101, VE-ID:80
VC Details: VC-ID:1080
Local MTU:1500, Remote MTU:1500
Remote (LB:25024,VBO:1,VBS:64) Local (LB:25024,VBO:65,VBS:64)
LB sent on known VEID:Yes
In Label:25039, Out Label:25033
PW Status:Established
VC Installed:Yes
VC Signaled Time: 00:02:02
Extended-Community Color:0

```

```
PE1#
```

PE2

```

PE2#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 101 | 30.0.1.1 | 24320 | 25033 | xe4 | 25039 | 2/Up | 1 B |
| GP | Active | 00:01:30 | - | | | | |

```

PE2#show mpls vpls de
debug-detail detail
PE2#show mpls vpls detail
Virtual Private LAN Service Instance: PE2-PE1, ID: 101
SIG-Protocol: BGP
Route-Distinguisher :80000:101
Route-Target :80000:101
VE-ID :80
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.q

```

```

Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe10.101
Status: Up
Subinterface Match Criteria(s) :
dot1q 101

Mesh Peers:
30.0.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:01:33)

```

```
PE2#show bgp l2vpn vpls detail
```

```

VPLS ID: 101
VE-ID: 80
Discovered Peers: 1
Route-Target: 80000:101
Local RD: 80000:101
All Local Label Blocks:
[LB:25024, VBO:1, VBS:64]
[LB:24960, VBO:65, VBS:64]
Mesh Peers:
BGP Peer:30.0.1.1/32
VC Nbr Address:30.0.1.1, RD:80000:101, VE-ID:10
VC Details: VC-ID:8010
Local MTU:1500, Remote MTU:1500
Remote (LB:25024,VBO:65,VBS:64) Local (LB:25024,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25033, Out Label:25039
PW Status:Established
VC Installed:Yes
VC Signaled Time: 00:01:38
Extended-Community Color:0

```

```
PE2#
```

End to end ping for VPLS service should be successful.

PE1

```

PE1#ping mpls vpls 101 peer 30.0.1.2/32 detail
Sending 5 MPLS Echos to VPLS Id : 101, timeout is 5 seconds

```

```

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

```

```
Type 'Ctrl+C' to abort
```

```

! seq_num = 1 10.0.1.12 1.90 ms
! seq_num = 2 10.0.1.12 0.61 ms
! seq_num = 3 10.0.1.12 0.53 ms
! seq_num = 4 10.0.1.12 0.53 ms
! seq_num = 5 10.0.1.12 0.57 ms

```

```

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.53/1.21/1.90
PE1#

```


PE2

```
PE2#ping mpls vpls 101 peer 30.0.1.1/32 detail
Sending 5 MPLS Echos to VPLS Id : 101, timeout is 5 seconds
```

Codes:

```
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed
```

Type 'Ctrl+C' to abort

```
! seq_num = 1 10.0.1.20 1.37 ms
! seq_num = 2 10.0.1.20 0.76 ms
! seq_num = 3 10.0.1.20 0.73 ms
! seq_num = 4 10.0.1.20 0.69 ms
! seq_num = 5 10.0.1.20 0.69 ms
```

```
Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.69/1.03/1.37
PE2#
```

Static VPLS Service Mapping Configuration

Overview


This chapter includes step-by-step configurations for static VPLS. It also contains an overview of the concepts of Static VPLS. Virtual Private LAN Service (VPLS) is a way to provide Ethernet-based multipoint-to-multipoint communication over IP- MPLS networks. It allows geographically-dispersed sites to share an Ethernet broadcast domain by connecting sites through pseudowires.

Configuration

Topology

The diagram depicts the topology for the configuration examples that follow.

Figure 77. Static VPLS service mapping



PE1: Loopback Interface

| | |
|---|---|
| <pre>#configure terminal</pre> | Enter configuration mode. |
| <pre>(config)#interface lo</pre> | Enter interface mode for the loopback interface. |
| <pre>(config-if)#ip address 1.1.1.1/32 secondary</pre> | Configure IP address on loopback interface. |
| <pre>(config-if)#commit</pre> | Commit the candidate configuration to the running configuration. |
| <pre>(config-if)#exit</pre> | Exit interface mode |
| <pre>(config)#router ldp</pre> | Enter router mode for LDP. |
| <pre>(config-router)#router-id 1.1.1.1</pre> | Set the router ID to IP address 1.1.1.1. |
| <pre>(config-router)#transport-address ipv4 1.1.1.1</pre> | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. <div><div></div><div>Note: It is preferable to use the loopback address as the transport address.</div></div> |
| <pre>(config-router)#targeted-peer ipv4 2.2.2.2</pre> | Configure targeted peer. |

| | |
|------------------------------------|-------------------------|
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |

PE1: Interface Configuration

| | |
|------------------------------------|--|
| (config)#interface ce2/1 | Enter interface mode for ce6/1. |
| (config-if)#ip address 10.1.1.1/24 | Configure IP address on the interface. |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on ce6/1. |
| (config-if)#commit | Commit the candidate configuration to the running configuration. |
| (config-if)#exit | Exit interface mode |

PE1: Static VPLS Configuration

| | |
|--|--|
| (config)#mpls vpls v251 251 | Enter VPLS configuration mode |
| (config-vpls)#vpls-peer 2.2.2.2 tunnel-id 1 manual | Configure VPLS peer |
| (config-vpls)#exit | Exit from VPLS configuration mode |
| (config)#mpls vpls v252 252 | Enter VPLS configuration mode |
| (config-vpls)#vpls-peer 2.2.2.2 tunnel-id 1 manual | Configure VPLS peer |
| (config-vpls)#exit | Exit from VPLS configuration mode |
| (config)#mpls vpls v253 253 | Enter VPLS configuration mode |
| (config-vpls)#vpls-peer 2.2.2.2 tunnel-id 1 manual | Configure VPLS peer |
| (config-vpls)#exit | Exit from VPLS configuration and configuration mode |
| (config)#mpls vpls v254 254 | Enter VPLS configuration mode |
| (config-vpls)#vpls-peer 2.2.2.2 tunnel-id 1 manual | Configure VPLS peer |
| (config-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-vpls)#end | Exit from VPLS configuration and configuration mode |


PE1: OSPF Configuration

| | |
|--|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-if)#ospf router-id 1.1.1.1 | Configure the router-id |
| (config-if)#network 1.1.1.1/32 area 0.0.0.0 | Advertise the loopback IP |
| (config-if)#network 10.1.1.0/32 area 0.0.0.0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the Interface address. |
| (config-if)#exit | Exit interface mode |

PE1: FIB Entry Configuration

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#mpls ftn-entry tunnel-id 1 2.2.2.2/32 100 10.1.1.2 ce2/1 primary | Configure Static LSP FTN entry |
| (config)#mpls ilm-entry 250 pop | Configure ILM entry |
| (config)#commit | Commit the candidate configuration to the running configuration. |
| (config)#exit | Exit |
| (config)#vpls fib-entry 251 peer 2.2.2.2 1001 xe11 2001 | Configuring VPLS FIB entry for vpls peer PE1. |
| (config)#vpls fib-entry 252 peer 2.2.2.2 1002 xe11 2002 | Configuring VPLS FIB entry for vpls peer PE1. |
| (config)#vpls fib-entry 253 peer 2.2.2.2 1003 xe11 2003 | Configuring VPLS FIB entry for vpls peer PE1. |
| (config)#vpls fib-entry 254 peer 2.2.2.2 1004 xe11 2004 | Configuring VPLS FIB entry for vpls peer PE1. |
| (config)#commit | Commit the candidate configuration to the running configuration. |

P: Loopback Interface

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter interface mode for the loopback interface. |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure IP address on loopback interface. |
| (config-if)#commit | Commit the candidate configuration to the running configuration. |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 4.4.4.4 | Set the router ID to IP address 4.4.4.4 |
| (config-router)#transport-address ipv4 4.4.4.4 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| <div>  Note: It is preferable to use the loopback address as the transport address. </div> | |
| (config-router)#exit | Exit router mode |

P: Interface Configuration

| | |
|-------------------------|--------------------------------|
| (config)#interface xe11 | Enter interface mode for xe11. |
|-------------------------|--------------------------------|

| | |
|------------------------------------|--|
| (config-if)#ip address 10.1.1.2/24 | Configure IP address on the interface. |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#commit | Commit the candidate configuration to the running configuration. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe13 | Enter interface mode for xe13. |
| (config-if)#ip address 20.1.1.2/24 | Configure IP address on the interface. |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#commit | Commit the candidate configuration to the running configuration. |
| (config-if)#exit | Exit interface mode |

P: FIB Entry Configuration

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#mpls ilm-entry 100 swap 200 xe13 20.1.1.2 2.2.2.2/32 | Configure Static LSP ILM entry |
| (config)#mpls ilm-entry 150 swap 250 xe11 10.1.1.2 1.1.1.1/32 | Configure ILM entry |
| (config)#commit | Commit the candidate configuration to the running configuration. |
| (config)#exit | Exit |

PE2: Loopback Interface

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter interface mode for the loopback interface. |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address on loopback interface. |
| (config-if)#commit | Commit the candidate configuration to the running configuration. |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| (config-router)#transport-address ipv4 2.2.2.2 | <p>Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run.</p> <div data-bbox="894 1705 967 1772" data-label="Image"> </div> <p>Note: It is preferable to use the loopback address as the transport address.</p> |

| | |
|--|--------------------------|
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |

PE2: Interface Configuration

| | |
|------------------------------------|--|
| (config)#interface xe15 | Enter interface mode for xe15. |
| (config-if)#ip address 20.1.1.1/24 | Configure IP address on the interface. |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#commit | Commit the candidate configuration to the running configuration. |
| (config-if)#exit | Exit interface mode |

OSPF Configuration

| | |
|--|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-if)#ospf router-id 2.2.2.2 | Configure the router-id |
| (config-if)#network 2.2.2.2/32 area 0.0.0.0 | Advertise the loopback ip |
| (config-if)#network 20.1.1.0/24 area 0.0.0.0 | Define the interface address on which the OSPF runs and associate an area ID(0) with the Interface address. |
| (config-if)#exit | Exit interface mode |

PE2: Static VPLS Configuration

| | |
|--|--|
| (config)#mpls vpls v251 251 | Enter VPLS configuration mode |
| (config-vpls)#vpls-peer 1.1.1.1 tunnel-id 1 manual | Configure VPLS peer |
| (config-vpls)#exit | Exit from VPLS configuration mode |
| (config)#mpls vpls v252 252 | Enter VPLS configuration mode |
| (config-vpls)#vpls-peer 1.1.1.1 tunnel-id 1 manual | Configure VPLS peer |
| (config-vpls)#exit | Exit from VPLS configuration mode |
| (config)#mpls vpls v253 253 | Enter VPLS configuration mode |
| (config-vpls)#vpls-peer 1.1.1.1 tunnel-id 1 manual | Configure VPLS peer |
| (config-vpls)#exit | Exit from VPLS configuration and configuration mode |
| (config)#mpls vpls v254 254 | Enter VPLS configuration mode |
| (config-vpls)#vpls-peer 1.1.1.1 tunnel-id 1 manual | Configure VPLS peer |
| (config-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-vpls)#exit | Exit from VPLS configuration and configuration mode |

PE2: FIB Entry Configuration

| | |
|--|--|
| (config)#mpls ftn-entry tunnel-id 1 1.1.1.1/32 150 20.1.1.2 xe15 | Configure Static LSP FTN entry |
| (config)#mpls ilm-entry 200 pop | Configure ILM entry |
| (config)#commit | Commit the candidate configuration to the running configuration. |
| (config)#exit | Exit |
| (config)#vpls fib-entry 251 peer 2.2.2.2 1001 xe11 2001 | Configuring VPLS FIB entry for vpls peer PE2. |
| (config)#vpls fib-entry 252 peer 2.2.2.2 1002 xe11 2002 | Configuring VPLS FIB entry for vpls peer PE2. |
| (config)#vpls fib-entry 253 peer 2.2.2.2 1003 xe11 2003 | Configuring VPLS FIB entry for vpls peer PE2. |
| (config)#vpls fib-entry 254 peer 2.2.2.2 1004 xe11 2004 | Configuring VPLS FIB entry for vpls peer PE2. |
| (config)#commit | Commit the candidate configuration to the running configuration. |

Configuration of SubInterface

PE1: Access port Configuration

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config)# interface xe2.254 switchport | Enter sub interface mode. |
| (config-if)# encapsulation dot1q 254 | Configure encapsulation type and value. |
| (config-acc-if-vpls)#access-if-vpls | Access VPLS under sub interface. |
| (config-acc-if-vpls)#mpls-vpls v254 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE1: POP

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)# interface xe0.251 switchport | Enter sub interface mode. |
| (config-if)# encapsulation dot1ad 3002 inner-dot1q 3005 | Configure encapsulation type and value. |
| (config-if)# rewrite pop | Configure rewrite with action pop. |
| (config-if)#access-if-vpls | Access VPLS under sub interface. |
| (config-acc-if-vpls)#mpls-vpls v251 | Associating the VPLS Instance to the attachment circuit interface. |

| | |
|-----------------------------|--|
| (config-acc-if-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE1: XLATE

| | |
|---|--|
| (config)# interface xe2.252 switchport | Enter sub interface mode. |
| (config-if)# encapsulation dot1ad 3110 inner-dot1q 3115 | Configure encapsulation type and value. |
| (config-if)# rewrite translate 2-to-1 dot1ad 3112 | Configure VLAN tag translation. |
| (config-if)#access-if-vpls | Access VPLS under sub interface. |
| (config-acc-if-vpls)#mpls-vpls v252 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE1: PUSH

| | |
|---|--|
| (config)# interface xe2.253 switchport | Enter sub interface mode. |
| (config-if)# encapsulation dot1ad 219 inner-dot1q 220 | Configure encapsulation type and value. |
| (config-if)# rewrite push dot1ad 221 | Configure rewrite push. |
| (config-acc-if-vpls)# access-if-vpls | Access VPLS under sub interface. |
| (config-acc-if-vpls)#mpls-vpls v253 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE2: Access port Configuration

| | |
|--|--|
| (config)# interface xe0.254 switchport | Enter sub interface mode. |
| (config-if)# encapsulation dot1q 254 | Configure encapsulation type and value. |
| (config-acc-if-vpls)#access-if-vpls | Access VPLS under sub interface. |
| (config-acc-if-vpls)#mpls-vpls v254 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-if-vpls)#exit | Exit from access mode. |

PE2: POP

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)# interface xe0.251 switchport | Enter sub interface mode. |
| (config-if)# encapsulation dot1ad 3002 inner-dot1q 3005 | Configure encapsulation type and value. |
| (config-if)# rewrite pop | Configure rewrite with action pop. |
| (config-if)#access-if-vpls | Access VPLS under sub interface. |
| (config-acc-if-vpls)#mpls-vpls v251 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE2: XLATE

| | |
|---|--|
| (config)# interface xe0.252 switchport | Enter sub interface mode. |
| (config-if)# encapsulation dot1ad 3110 inner-dot1q 3115 | Configure encapsulation type and value. |
| (config-if)# rewrite translate 2-to-1 dot1ad 3112 | Configure VLAN tag translation. |
| (config-if)#access-if-vpls | Access VPLS under sub interface. |
| (config-acc-if-vpls)#mpls-vpls v252 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE2: PUSH

| | |
|---|--|
| (config)# interface xe0.253 switchport | Enter sub interface mode. |
| (config-if)# encapsulation dot1ad 219 inner-dot1q 220 | Configure encapsulation type and value. |
| (config-if)# rewrite push dot1ad 221 | Configure rewrite push. |
| (config-acc-if-vpls)# access-if-vpls | Access VPLS under sub interface. |
| (config-acc-if-vpls)#mpls-vpls v253 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#commit | Commit the candidate configuration to the running configuration. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

Validation

PE1

```
PE1#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 251 | 2.2.2.2 | 151 | 1001 | N/A | 2001 | 2/Up | 3 |
| TATIC | Active | 00:16:45 | - | | | | S |
| 252 | 2.2.2.2 | 151 | 1002 | N/A | 2002 | 2/Up | 4 |
| TATIC | Active | 00:15:57 | - | | | | S |
| 253 | 2.2.2.2 | 151 | 1003 | N/A | 2003 | 2/Up | 5 |
| TATIC | Active | 00:15:57 | - | | | | S |
| 254 | 2.2.2.2 | 151 | 1004 | N/A | 2004 | 2/Up | 6 |
| TATIC | Active | 00:15:57 | - | | | | S |

```
PE1#sh mpls vpls detail
Virtual Private LAN Service Instance: v251, ID: 251
SIG-Protocol: STATIC
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:
```

```
Configured interfaces:
Interface: xe2.251
Status: Up
Subinterface Match Criteria(s) :
dot1q 251
```

```
Mesh Peers:
2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:01:16)
Tunnel-Id: 1
```

```
Virtual Private LAN Service Instance: v252, ID: 252
SIG-Protocol: STATIC
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:
```

```
Configured interfaces:
Interface: xe2.252
Status: Up
Subinterface Match Criteria(s) :
dot1q 3110 inner-dot1q 3115
```

```
Mesh Peers:
2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:01:16)
Tunnel-Id: 1
```

```

Virtual Private LAN Service Instance: v253, ID: 253
  SIG-Protocol: STATIC
  Attachment-Circuit: UP
  Learning: Enabled
  Control-Word: Disabled
  Flow Label Status: Disabled, Direction: None, Static: No
  Group ID: 0, Configured MTU: 1500
  Description: none
  service-tpid: dot1q
  Operating mode: Raw
  MAC Withdrawal:

Configured interfaces:
  Interface: xe2.253
  Status: Up
  Subinterface Match Criteria(s) :
  dot1ad 219 inner-dot1q 220

Mesh Peers:
  2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:01:16)
  Tunnel-Id: 1

Virtual Private LAN Service Instance: v254, ID: 254
  SIG-Protocol: STATIC
  Attachment-Circuit: UP
  Learning: Enabled
  Control-Word: Disabled
  Flow Label Status: Disabled, Direction: None, Static: No
  Group ID: 0, Configured MTU: 1500
  Description: none
  service-tpid: dot1q
  Operating mode: Raw
  MAC Withdrawal:

Configured interfaces:
  Interface: xe2.254
  Status: Up
  Subinterface Match Criteria(s) :
  dot1q 254

Mesh Peers:
  2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:01:16)
  Tunnel-Id: 1

```

PE2

```

PE2#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 251 | 1.1.1.1 | 501 | 1001 | N/A | 2001 | 2/Up | 4 S |
| TATIC | Active | 00:16:34 | - | | | | |
| 252 | 1.1.1.1 | 501 | 2002 | N/A | 1002 | 2/Up | 10 S |
| TATIC | Active | 00:04:38 | - | | | | |
| 253 | 1.1.1.1 | 501 | 2003 | N/A | 1003 | 2/Up | 12 S |
| TATIC | Active | 00:14:48 | - | | | | |
| 254 | 1.1.1.1 | 501 | 2004 | N/A | 1004 | 2/Up | 13 S |
| TATIC | Active | 00:14:48 | - | | | | |

```

PE2#sh mpls vpls detail
Virtual Private LAN Service Instance: v251, ID: 251
  SIG-Protocol: STATIC

```

```
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe0.251
Status: Up
Subinterface Match Criteria(s) :
dot1q 251

Mesh Peers:
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:02:01)
Tunnel-Id: 1

Virtual Private LAN Service Instance: v252, ID: 252
SIG-Protocol: STATIC
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe0.252
Status: Up
Subinterface Match Criteria(s) :
dot1q 3110 inner-dot1q 3115

Mesh Peers:
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:02:01)
Tunnel-Id: 1

Virtual Private LAN Service Instance: v253, ID: 253
SIG-Protocol: STATIC
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe0.253
Status: Up
Subinterface Match Criteria(s) :
dot1ad 219 inner-dot1q 220

Mesh Peers:
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:02:01)
Tunnel-Id: 1

Virtual Private LAN Service Instance: v254, ID: 254
```

```
SIG-Protocol: STATIC
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
  Interface: xe0.254
  Status: Up
  Subinterface Match Criteria(s) :
    dot1q 254

Mesh Peers:
  1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:02:01)
  Tunnel-Id: 1
```

LDP-VPLS Service Mapping Configuration


Overview

This chapter includes step-by-step configurations for LDP VPLS. It also contains an overview of the concepts of LDP VPLS. Virtual Private LAN Service (VPLS) is a way to provide Ethernet-based multipoint-to-multipoint communication over IP- MPLS networks. It allows geographically-dispersed sites to share an Ethernet broadcast domain by connecting sites through pseudowires.

Topology

The diagram depicts the topology for the configuration examples that follow.

Figure 78. LDP-VPLS service mapping



Configuration

PE1: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure IP address on loopback interface. |
| (config-if)#exit | Exit interface mode |

PE1: Global LDP

| | |
|--|---|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1 |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure transport address |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Configure targeted peer |
| (config-router)#targeted-peer #commit | Commit candidate configuration to the running configuration |
| (config-router-targeted-peer)#end | Exit from router target peer and LDP mode |

PE1: Interface Configuration

| | |
|-------------------------------------|--|
| #configure terminal | Enter configuration mode |
| (config)#interface ce2/1 | Enter the Interface mode for ce2/1. |
| (config-if)# ip address 10.1.1.1/24 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |

PE1: OSPF Configuration

| | |
|--|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 1.1.1.1 | Router-id configurations |
| (config-router)#network 1.1.1.1/32 area 0 | Advertise loopback address in OSPF. |
| (config-router)#network 10.1.1.0/24 area 0 | Advertise network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |

PE1: LDP VPLS Configuration

| | |
|-------------------------------------|----------------------------------|
| (config)#mpls vpls v1 25 | Enter VPLS config mode |
| (config-vpls)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 2.2.2.2 | Configure VPLS Peer |
| (config-vpls-sig)#exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v2 26 | Enter VPLS config mode |
| (config-vpls)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 2.2.2.2 | Configure VPLS Peer |
| (config-vpls-sig)#exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v3 27 | Enter VPLS config mode |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 2.2.2.2 | Configure VPLS Peer |

| | |
|--------------------------------------|---|
| (config-vpls-sig) # exit-signaling | Exit Signaling LDP mode |
| (config-vpls) #exit | Exit VPLS mode |
| (config) #mpls vpls v4 28 | Enter VPLS config mode |
| (config-vpls) #signaling ldp | Define Signaling as LDP |
| (config-vpls-sig) #vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig) #vpls-peer 2.2.2.2 | Configure VPLS Peer |
| (config-vpls-sig) # exit-signaling | Exit Signaling LDP mode |
| (config-vpls) #exit | Exit VPLS mode |
| (config) #commit | Commit candidate configuration to the running configuration |

P: Loopback Interface

| | |
|--|--|
| #configure terminal | Enter configuration mode. |
| (config) #interface lo | Enter the Interface mode for the loopback interface. |
| (config-if) #ip address 4.4.4.4/32 secondary | Configure IP address on loopback interface. |
| (config-if) #exit | Exit interface mode |

P: Global LDP

| | |
|---|---|
| (config) #router ldp | Enter the Router LDP mode. |
| (config-router) #router-id 4.4.4.4 | Set the router ID to IP address 4.4.4.4 |
| (config-router) #transport-address ipv4 4.4.4.4 | Configure transport address |
| (config-router) #commit | Commit candidate configuration to the running configuration |
| (config-router-targeted-peer) #end | Exit from router target peer and LDP mode |

P: Interface Configuration

| | |
|--------------------------------------|--|
| #configure terminal | Enter configuration mode |
| (config) #interface xe11 | Enter the Interface mode for xe11. |
| (config-if) # ip address 10.1.1.2/24 | Configure IP address on the interface. |
| (config-if) #enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if) # label-switching | Enable label switching on the interface. |
| (config-if) #exit | Exit interface mode |
| (config) #interface xe13 | Enter the Interface mode for xe13. |
| (config-if) # ip address 20.1.1.2/24 | Configure IP address on the interface. |
| (config-if) #enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if) # label-switching | Enable label switching on the interface. |
| (config-if) #exit | Exit interface mode |

P: OSPF Configuration

| | |
|--|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#network 4.4.4.4/32 area 0 | Advertise loopback address in OSPF. |
| (config-router)#network 20.1.1.0/24 area 0 | Advertise network address in OSPF. |
| (config-router)#network 10.1.1.0/24 area 0 | Advertise network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit candidate configuration to the running configuration |

PE2: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure IP address on loopback interface. |
| (config-if)#exit | Exit interface mode |

PE2: Global LDP

| | |
|--|---|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure transport address |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer |
| (config-router-targeted-peer)#commit | Commit candidate configuration to the running configuration |
| (config-router-targeted-peer)#end | Exit from router target peer and LDP mode |

PE2: Interface Configuration

| | |
|-------------------------------------|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface xe15 | Enter the Interface mode for xe1. |
| (config-if)# ip address 20.1.1.1/24 | Configure IP address on the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |

PE2: OSPF Configuration

| | |
|---|-------------------------------------|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#network 2.2.2.2/32 area 0 | Advertise loopback address in OSPF. |

| | |
|--|---|
| (config-router)#network 20.1.1.0/24 area 0 | Advertise network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |

PE2: LDP VPLS Configuration

| | |
|-------------------------------------|---|
| (config)#mpls vpls v1 25 | Enter VPLS config mode |
| (config-vpls)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 1.1.1.1 | Configure VPLS Peer |
| (config-vpls-sig)# exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v2 26 | Enter VPLS config mode |
| (config-vpls)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 1.1.1.1 | Configure VPLS Peer |
| (config-vpls-sig)# exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v3 27 | Enter VPLS config mode |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 1.1.1.1 | Configure VPLS Peer |
| (config-vpls-sig)# exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v4 28 | Enter VPLS config mode |
| (config-vpls)#signaling ldp | Define Signaling as LDP |
| (config-vpls-sig)#vpls-type vlan | Type VLAN configuration for VPLS |
| (config-vpls-sig)#vpls-peer 1.1.1.1 | Configure VPLS Peer |
| (config-vpls-sig)# exit-signaling | Exit Signaling LDP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#commit | Commit candidate configuration to the running configuration |

LDP VPLS Subinterface Matching

PE1: POP

| | |
|--|---|
| #configure terminal | Configure mode |
| (config)# int xe2.25 switchport | Configure the interface xe2.25 as a Layer 2 switch port. |
| (config-if)#encapsulation dot1ad 500 inner-dot1q 501 | Configure encapsulation to dot1ad with outer VLAN ID 500 and inner VLAN ID 501. |
| (config-if)#rewrite pop | Configure rewrite with action pop. |
| (config-if)#access-if-vpls | Entering Access interface mode. |
| (config-acc-if-vpls)#mpls-vpls v6 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exiting out of access interface mode. |

PE1: XLATE

| | |
|---|---|
| (config)# int xe2.26 switchport | Configure the interface xe2.26 as a Layer 2 switch port. |
| (config-if)# encapsulation dot1ad 500 inner-dot1q 501 | Configure encapsulation to dot1ad with outer VLAN ID 500 and inner VLAN ID 501. |
| (config-if)#rewrite translate 2-to-1 dot1q 250 | Configure VLAN tag (250) translation on the interface. |
| (config-if)#access-if-vpls | Entering Access interface mode. |
| (config-acc-if-vpls)mpls-vpls v2 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls))#exit | Exiting out of access interface mode. |

PE1: PUSH

| | |
|--|---|
| (config)# interface xe2.27 switchport | Configure the interface xe2.27 as a Layer 2 switch port. |
| (config-if)#encapsulation dot1ad 251 inner-dot1q 261 | Configure encapsulation to dot1ad with outer VLAN ID 251 and inner VLAN ID 261. |
| (config-if)#rewrite push dot1q 1550 | Configure interface to push a single VLAN tag with VLAN ID 1550. |
| (config-if)#access-if-vpls | Entering Access interface mode. |
| (config-acc-if-vpls)#mpls-vpls v3 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exiting out of access interface mode. |

PE1:POP-2tag

| | |
|---|---|
| (config)# interface xe2.28 switchport | Configure the interface xe2.28 as a Layer 2 switch port. |
| (config-if)# encapsulation dot1ad 400 inner-dot1q 401 | Configure encapsulation to dot1ad with outer VLAN ID 400 and inner VLAN ID 401. |
| (config-if)# rewrite pop-2tag | Configure rewrite with action pop-2tag. |
| (config-if)#access-if-vpls | Entering Access interface mode. |
| (config-acc-if-vpls)#mpls-vpls v4 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exiting out of access interface mode. |

PE2: POP

| | |
|--|---|
| #configure terminal | Configure mode |
| (config)# int xe0.25 switchport | Configure the interface xe0.25 as a Layer 2 switch port. |
| (config-if)#encapsulation dot1ad 500 inner-dot1q 501 | Configure encapsulation to dot1ad with outer VLAN ID 500 and inner VLAN ID 501. |
| (config-if)#rewrite pop | Configure rewrite with action pop. |
| (config-if)#access-if-vpls | Entering Access interface mode. |
| (config-acc-if-vpls)mpls-vpls v1 | Associating the VPLS Instance to the attachment circuit interface. |
| ((config-acc-if-vpls))#exit | Exiting out of access interface mode. |

PE2: XLATE

| | |
|---|---|
| (config)# int xe0.26 switchport | Configure the interface xe0.26 as a Layer 2 switch port. |
| (config-if)# encapsulation dot1ad 500 inner-dot1q 501 | Configure encapsulation to dot1ad with outer VLAN ID 500 and inner VLAN ID 501. |
| (config-if)#rewrite translate 2-to-1 dot1q 250 | Configure rewrite VLAN tag translation. |
| (config-if)#access-if-vpls | Entering Access interface mode. |
| (config-acc-if-vpls)mpls-vpls v2 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exiting out of access interface mode. |

PE2: PUSH

| | |
|--|--|
| (config)# interface xe0.27 switchport | Configure the interface xe0.27 as a Layer 2 switch port. |
| (config-if)#encapsulation dot1ad 251 inner-dot1q 261 | Configure encapsulation to dot1ad with outer |

| | |
|-------------------------------------|--|
| | VLAN ID 251 and inner VLAN ID 261. |
| (config-if)#rewrite push dot1q 1550 | Configure rewrite VLAN tag translation. |
| (config-if)#access-if-vpls | Entering Access interface mode. |
| (config-acc-if-vpls)#mpls-vpls v3 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exiting out of access interface mode. |

PE2:POP-2tag

| | |
|---|---|
| (config)# interface xe0.28 switchport | Configure the interface xe0.28 as a Layer 2 switch port. |
| (config-if)# encapsulation dot1ad 400 inner-dot1q 401 | Configure encapsulation to dot1ad with outer VLAN ID 400 and inner VLAN ID 401. |
| (config-if)# rewrite pop-2tag | Configure rewrite with action pop-2tag. |
| (config-if)#access-if-vpls | Entering Access mode for EVPN MPLS ID configuration. |
| (config-acc-if-vpls)#mpls-vpls v4 | Associating the VPLS Instance to the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exiting out of access interface mode. |

Validation

PE1

```
PE1#sh mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | SIG- |
|----------|-----------|--------------|-----------|--------------|-----------|---------|----------|------|
| Protocol | Status | UpTime | Ext-Color | | | | | |
| 25 | 2.2.2.2 | 120 | 24960 | ce2/1 | 24961 | 2/Up | 3 | LDP |
| | Active | 00:02:00 | - | | | | | |
| 26 | 2.2.2.2 | 120 | 24962 | ce2/1 | 24962 | 2/Up | 2 | LDP |
| | Active | 00:02:00 | - | | | | | |
| 27 | 2.2.2.2 | 120 | 24961 | ce2/1 | 24963 | 2/Up | 4 | LDP |
| | Active | 00:02:00 | - | | | | | |
| 28 | 2.2.2.2 | 120 | 24963 | ce2/1 | 24964 | 2/Up | 5 | LDP |
| | Active | 00:00:11 | - | | | | | |

```
PE1#sh ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
Session has to be cleared manually
```

| Code | Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|------|-----------------|---------|---------|-------------|-----------|----------|
| | 4.4.4.4 | ce2/1 | Passive | OPERATIONAL | 30 | 1d23h13m |
| | 2.2.2.2 | ce2/1 | Passive | OPERATIONAL | 30 | 07:27:27 |

```
PE1#sh ldp vpls detail
VPLS Identifier : 25
Peer IP : 2.2.2.2
VC State : UP
```

```
VC Type           : vlan
VC Label Sent     : 24960
VC Label Received : 24961
Local MTU         : 1500
Remote MTU        : 1500
Local Control Word : disabled
Remote Control Word : disabled
Current use       : disabled
Local Flow Label   : disabled
Remote Flow Label  : disabled
Current use       : disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
Local FEC Type    : 128
Remote FEC Type   : 128
LDP-VPLS Signaled Time : 00:04:14
```

```
VPLS Identifier   : 26
Peer IP           : 2.2.2.2
VC State          : UP
VC Type           : vlan
VC Label Sent     : 24962
VC Label Received : 24962
Local MTU         : 1500
Remote MTU        : 1500
Local Control Word : disabled
Remote Control Word : disabled
Current use       : disabled
Local Flow Label   : disabled
Remote Flow Label  : disabled
Current use       : disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
Local FEC Type    : 128
Remote FEC Type   : 128
LDP-VPLS Signaled Time : 00:04:14
```

```
VPLS Identifier   : 27
Peer IP           : 2.2.2.2
VC State          : UP
VC Type           : vlan
VC Label Sent     : 24961
VC Label Received : 24963
Local MTU         : 1500
Remote MTU        : 1500
Local Control Word : disabled
Remote Control Word : disabled
Current use       : disabled
Local Flow Label   : disabled
Remote Flow Label  : disabled
Current use       : disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
Local FEC Type    : 128
Remote FEC Type   : 128
LDP-VPLS Signaled Time : 00:04:14
```

```
VPLS Identifier   : 28
Peer IP           : 2.2.2.2
VC State          : UP
VC Type           : vlan
VC Label Sent     : 24963
VC Label Received : 24964
Local MTU         : 1500
Remote MTU        : 1500
```

```

Local Control Word : disabled
Remote Control Word: disabled
Current use       : disabled
Local Flow Label  : disabled
Remote Flow Label : disabled
Current use       : disabled
Local PW Status Capability : disabled
Remote PW Status Capability : disabled
Current PW Status TLV : disabled
Local FEC Type    : 128
Remote FEC Type   : 128
LDP-VPLS Signaled Time : 00:02:25

```

```
PE1#sh ldp vpls
```

| VPLS-ID | VC-Style | Peer-Address | State | Type | Label-Sent | Label-Rcvd | Local-PW- |
|---------|----------|------------------|-------|------|------------|------------|-----------|
| Status | | Remote-PW-Status | | | | | |
| 25 | Mesh | 2.2.2.2 | UP | vlan | 24960 | 24961 | N/A |
| | N/A | | | | | | |
| 26 | Mesh | 2.2.2.2 | UP | vlan | 24962 | 24962 | N/A |
| | N/A | | | | | | |
| 27 | Mesh | 2.2.2.2 | UP | vlan | 24961 | 24963 | N/A |
| | N/A | | | | | | |
| 28 | Mesh | 2.2.2.2 | UP | vlan | 24963 | 24964 | N/A |
| | N/A | | | | | | |

```
PE1#show mpls vpls detail
```

```

Virtual Private LAN Service Instance: v1, ID: 25
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 88a8
MAC Withdrawal:

```

```
Configured interfaces:
```

```

Interface: xe2.25
Status: Up
Subinterface Match Criteria(s) :
dot1ad 500 inner-dot1q 501

```

```
Mesh Peers:
```

```

2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:07:15)
FEC signaling element: FEC128

```

```
Virtual Private LAN Service Instance: v2, ID: 26
```

```

SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 88a8
MAC Withdrawal:

```

```
Configured interfaces:
```

```

Interface: xe2.26
Status: Up
Subinterface Match Criteria(s) :
dotlad 503 inner-dot1q 504

Mesh Peers:
  2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:07:15)
    FEC signaling element: FEC128

Virtual Private LAN Service Instance: v3, ID: 27
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 88a8
MAC Withdrawal:

Configured interfaces:
Interface: xe2.27
Status: Up
Subinterface Match Criteria(s) :
dotlad 251 inner-dot1q 261

Mesh Peers:
  2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:07:15)
    FEC signaling element: FEC128

Virtual Private LAN Service Instance: v4, ID: 28
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 88a8
MAC Withdrawal:

Configured interfaces:
Interface: xe2.28
Status: Up
Subinterface Match Criteria(s) :
dotlad 400 inner-dot1q 401

Mesh Peers:
  2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:05:26)
    FEC signaling element: FEC128

```

PE2

```

PE2#sh mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |

| | | | | | | | | | |
|----|---------|----------|---|-------|------|-------|------|---|---|
| 25 | 1.1.1.1 | 150 | | 24961 | xe15 | 24960 | 2/Up | 2 | L |
| DP | Active | 00:06:11 | - | | | | | | |
| 26 | 1.1.1.1 | 150 | | 24962 | xe15 | 24962 | 2/Up | 3 | L |
| DP | Active | 00:06:11 | - | | | | | | |
| 27 | 1.1.1.1 | 150 | | 24963 | xe15 | 24961 | 2/Up | 4 | L |
| DP | Active | 00:06:11 | - | | | | | | |
| 28 | 1.1.1.1 | 150 | | 24964 | xe15 | 24963 | 2/Up | 5 | L |
| DP | Active | 00:04:21 | - | | | | | | |

```
PE2#sh ldp vpls
```

| VPLS-ID | VC-Style | Peer-Address | State | Type | Label-Sent | Label-Rcvd | Local-PW-Status |
|---------|----------|--------------|-------|------|------------|------------|-----------------|
| 25 | Mesh | 1.1.1.1 | UP | vlan | 24966 | 24968 | Forwarding |
| 26 | Mesh | 1.1.1.1 | UP | vlan | 24967 | 24969 | Forwarding |
| 27 | Mesh | 1.1.1.1 | UP | vlan | 24968 | 24970 | Forwarding |
| 28 | Mesh | 1.1.1.1 | UP | vlan | 24971 | 24969 | Forwarding |

```
PE2#sh mpls vpls detail
```

```
Virtual Private LAN Service Instance: v1, ID: 25
```

```
SIG-Protocol: LDP
```

```
Attachment-Circuit: UP
```

```
Learning: Enabled
```

```
Control-Word: Disabled
```

```
Flow Label Status: Disabled, Direction: None, Static: No
```

```
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
```

```
Description: none
```

```
service-tpid: dot1.ad
```

```
Operating mode: Tagged
```

```
Svlan Id: 0
```

```
Svlan Tpid: 88a8
```

```
MAC Withdrawal:
```

```
Configured interfaces:
```

```
Interface: xe0.25
```

```
Status: Up
```

```
Subinterface Match Criteria(s) :
```

```
dot1ad 500 inner-dot1q 501
```

```
Mesh Peers:
```

```
1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:06:59)
```

```
FEC signaling element: FEC128
```

```
Virtual Private LAN Service Instance: v2, ID: 26
```

```
SIG-Protocol: LDP
```

```
Attachment-Circuit: UP
```

```
Learning: Enabled
```

```
Control-Word: Disabled
```

```
Flow Label Status: Disabled, Direction: None, Static: No
```

```
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
```

```
Description: none
```

```
service-tpid: dot1.ad
```

```
Operating mode: Tagged
```

```
Svlan Id: 0
```

```
Svlan Tpid: 88a8
```

```
MAC Withdrawal:
```

```
Configured interfaces:
```

```
Interface: xe0.26
```

```
Status: Up
```

```
Subinterface Match Criteria(s) :
```

```
dot1ad 503 inner-dot1q 504
```

Mesh Peers:

1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:06:59)
FEC signaling element: FEC128

Virtual Private LAN Service Instance: v3, ID: 27

SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 88a8
MAC Withdrawal:

Configured interfaces:

Interface: xe0.27
Status: Up
Subinterface Match Criteria(s) :
dotlad 251 inner-dot1q 261

Mesh Peers:

1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:06:59)
FEC signaling element: FEC128

Virtual Private LAN Service Instance: v4, ID: 28

SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 88a8
MAC Withdrawal:

Configured interfaces:

Interface: xe0.28
Status: Up
Subinterface Match Criteria(s) :
dotlad 400 inner-dot1q 401

Mesh Peers:

1.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:05:09)
FEC signaling element: FEC128

BGP-VPLS Service Mapping Configuration


Overview

This chapter provides step-by-step instructions to configure BGP Virtual Private LAN Service (VPLS), including an overview of the BGP VPLS concept. VPLS is a method that enables Ethernet-based multipoint-to-multipoint communication over IP-MPLS networks. It allows geographically separated sites to be part of the same Ethernet broadcast domain by linking them with pseudowires.

Topology

The diagram depicts the topology for the configuration examples that follow.

Figure 79. BGP-VPLS service mapping



Configuration

PE1: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 1.1.1.1/32 secondary | Configure the IP address on the loopback interface. |
| (config-if)#exit | Exit interface mode |

PE1: Interface Configuration

| | |
|-------------------------------------|--|
| (config)#interface ce2/1 | Enter the Interface mode for ce2/1. |
| (config-if)# ip address 10.1.1.1/24 | Configure the IP address on the interface. |
| (config-if)#enable-rsvp | Enable RSVP on the physical interface |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |

PE1: OSPF Configuration

| | |
|--|-----------------------------|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#ospf router-id 1.1.1.1 | Router-id configurations |

| | |
|--|---|
| (config-router)#network 1.1.1.1/32 area 0 | Advertise the loopback address in OSPF. |
| (config-router)#network 10.1.1.0/24 area 0 | Advertise the network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |

PE1: Global RSVP

| | |
|----------------------|---|
| (config)#router rsvp | Enter the Router OSPF mode. |
| (config-router)#exit | Exit Router RSVP mode and return to Configure mode. |

PE1: RSVP-Trunk Configuration

| | |
|---------------------------|--|
| (config)#rsvp-trunk 1 | Enter the Trunk configuration mode |
| (config-trunk)#to 2.2.2.2 | Configure the destination of the Trunk |
| (config-trunk)#exit | Exit.Trunk configuration mode |

PE1: BGP Configuration

| | |
|--|---|
| (config)# router bgp 100 | Enter the BGP configuration mode. |
| (config-router)#neighbor 2.2.2.2 remote-as 100 | Configure neighbor |
| (config-router)#neighbor 2.2.2.2 update-source 1.1.1.1 | Update the loopback address as the source |
| (config-router)#address-family l2vpn vpls | Enter address family mode. |
| (config-router-af)#neighbor 2.2.2.2 activate | Activate the neighbor. |
| (config-router-af)#exit | Exit the address family mode. |
| (config-router)#exit | Exit Router BGP mode |

PE1: BGP VPLS Configuration

| | |
|-------------------------------|------------------------------|
| (config)#mpls vpls v105 105 | Enter VPLS config mode |
| (config)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling bgp | Define Signaling as BGP |
| (config-vpls-sig)#ve-id 1 | Configure VE-ID |
| (config-vpls-sig)#exit | Exit Signaling BGP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v106 106 | Enter VPLS config mode |
| (config)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling bgp | Define Signaling as BGP |
| (config-vpls-sig)#ve-id 1 | Configure VE-ID |
| (config-vpls-sig)#exit | Exit Signaling BGP mode |

| | |
|-----------------------------|---|
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v107 107 | Enter VPLS config mode |
| (config-vpls)#signaling bgp | Define Signaling as BGP |
| (config-vpls-sig)#ve-id 1 | Configure VE-ID |
| (config-vpls-sig)#exit | Exit Signaling BGP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v108 108 | Enter VPLS config mode |
| (config-vpls)#signaling bgp | Define Signaling as BGP |
| (config-vpls-sig)#ve-id 1 | Configure VE-ID |
| (config-vpls-sig)#exit | Exit Signaling BGP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#commit | Commit the candidate configuration to the running configuration |

P1: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 4.4.4.4/32 secondary | Configure the IP address on the loopback interface. |
| (config-if)#exit | Exit interface mode |

P1: Interface Configuration

| | |
|------------------------------------|--|
| (config)#interface xe11 | Enter the Interface mode for xe11 |
| (config-if)#ip address 10.1.1.2/24 | Configure the IP address on the interface. |
| (config-if)#enable-rsvp | Enable RSVP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe13 | Enter the Interface mode for xe13 |
| (config-if)#ip address 20.1.1.2/24 | Configure the IP address on the interface. |
| (config-if)#enable-rsvp | Enable RSVP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |

P1: OSPF Configuration

| | |
|--|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#network 4.4.4.4/32 area 0 | Advertise the loopback address in OSPF. |
| (config-router)#network 10.1.1.0/24 area 0 | Advertise the network address in OSPF. |

| | |
|--|---|
| (config-router)#network 20.1.1.0/24 area 0 | Advertise the network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |

P1: Global RSVP

| | |
|----------------------|---|
| (config)#router rsvp | Enter the Router OSPF mode. |
| (config-router)#exit | Exit Router RSVP mode and return to Configure mode. |
| (config)#commit | Commit the candidate configuration to the running configuration |

PE2: Loopback Interface

| | |
|---|--|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)#ip address 2.2.2.2/32 secondary | Configure the IP address on the loopback interface. |
| (config-if)#exit | Exit interface mode |

PE2: Interface Configuration

| | |
|------------------------------------|--|
| (config)#interface xe15 | Enter the Interface mode for xe15 |
| (config-if)#ip address 20.1.1.1/24 | Configure the IP address on the interface. |
| (config-if)#enable-rsvp | Enable RSVP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface. |
| (config-if)#exit | Exit interface mode |

PE2: OSPF Configuration

| | |
|--|---|
| (config)#router ospf 100 | Enter the Router OSPF mode. |
| (config-router)#network 2.2.2.2/32 area 0 | Advertise the loopback address in OSPF. |
| (config-router)#network 20.1.1.0/24 area 0 | Advertise the network address in OSPF. |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |

PE2: Global RSVP

| | |
|----------------------|---|
| (config)#router rsvp | Enter the Router OSPF mode. |
| (config-router)#exit | Exit Router RSVP mode and return to Configure mode. |

PE2: RSVP-Trunk Configuration

| | |
|-----------------------|------------------------------------|
| (config)#rsvp-trunk 1 | Enter the Trunk configuration mode |
|-----------------------|------------------------------------|

| | |
|---------------------------|--|
| (config-trunk)#to 1.1.1.1 | Configure the destination of the Trunk |
| (config-trunk)#exit | Exit.Trunk configuration mode |

PE2: BGP Configuration

| | |
|--|---|
| (config)# router bgp 100 | Enter the BGP configuration mode. |
| (config-router)#neighbor 1.1.1.1 remote-as 100 | Configure neighbor |
| (config-router)#neighbor 1.1.1.1 update-source 2.2.2.2 | Update the loopback address as the source |
| (config-router)#address-family l2vpn vpls | Enter address family mode. |
| (config-router-af)#neighbor 1.1.1.1 activate | Activate the neighbor. |
| (config-router-af)#exit | Exit the address family mode. |
| (config-router)#exit | Exit Router BGP mode |

PE2: BGP VPLS Configuration

| | |
|-------------------------------|---|
| (config)#mpls vpls v105 105 | Enter VPLS config mode |
| (config)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling bgp | Define Signaling as BGP |
| (config-vpls-sig)#ve-id 2 | Configure VE-ID |
| (config-vpls-sig)#exit | Exit Signaling BGP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v106 106 | Enter VPLS config mode |
| (config)#service-tpid dot1.ad | Service tp-id configuration. |
| (config-vpls)#signaling bgp | Define Signaling as BGP |
| (config-vpls-sig)#ve-id 2 | Configure VE-ID |
| (config-vpls-sig)#exit | Exit Signaling BGP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v107 107 | Enter VPLS config mode |
| (config-vpls)#signaling bgp | Define Signaling as BGP |
| (config-vpls-sig)#ve-id 2 | Configure VE-ID |
| (config-vpls-sig)#exit | Exit Signaling BGP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#mpls vpls v108 108 | Enter VPLS config mode |
| (config-vpls)#signaling bgp | Define Signaling as BGP |
| (config-vpls-sig)#ve-id 2 | Configure VE-ID |
| (config-vpls-sig)#exit | Exit Signaling BGP mode |
| (config-vpls)#exit | Exit VPLS mode |
| (config)#commit | Commit the candidate configuration to the running configuration |

BGP VPLS SUBINTERFACE ATTACHMENT

PE1: POP

| | |
|---|---|
| (config)#interface xe2.105 switchport | Create an L2 sub-interface xe2.105 on the physical interface xe2. |
| (config-if)# encapsulation dot1ad 3001 inner-dot1q 2001 | Configure encapsulation using the outer VLAN ID 3001 (dot1ad) and the inner VLAN ID 2001 (dot1q). |
| (config-if)# rewrite pop | Configure rewrite with action pop. |
| (config-if)# access-if-vpls | Access VPLS under the sub-interface. |
| (config-acc-if-vpls)# mpls-vpls v105 | Associate the VPLS instance with the attachment circuit interface. |
| (config -acc-if-vpls)#exit | Exit from access mode. |

PE1: XLATE

| | |
|---|--|
| (config)# interface xe2.106 switchport | Create L2 sub-interface xe2.106 on the physical interface xe2. |
| (config-if)# encapsulation dot1ad 510 inner-dot1q 511 | Configure encapsulation with the outer VLAN ID 510 (dot1ad) and the inner VLAN ID 511 (dot1q). |
| (config-if)# rewrite translate 2-to-1 dot1q 251 | Configure VLAN tag (251) translation on the interface. |
| (config-if)# access-if-vpls | Access VPLS under the sub-interface. |
| (config-acc-if-vpls)# mpls-vpls v106 | Associate the VPLS instance with the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE1: PUSH

| | |
|---|--|
| (config)#interface xe2.107 switchport | Create L2 sub-interface xe2.107 on the physical interface xe2. |
| (config-if)# encapsulation dot1ad 301 inner-dot1q 310 | Configure encapsulation with the outer VLAN ID 301 (dot1ad) and the inner VLAN ID 310 (dot1q). |
| (config-if)# rewrite push dot1q 321 | Configure the interface to push a single VLAN tag with VLAN ID 321. |
| (config-if)# access-if-vpls | Access VPLS under the sub-interface. |
| (config-acc-if-vpls)# mpls-vpls v107 | Associate the VPLS instance with the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE1: POP-2TAG

| | |
|---------------------------------------|---|
| (config)#interface xe2.108 switchport | Create L2 sub-interface xe2.108 on the physical |
|---------------------------------------|---|

| | |
|---|--|
| | interface xe2. |
| (config-if)# encapsulation dot1ad 105 inner-dot1q 115 | Configure encapsulation with the outer VLAN ID 105 (dot1ad) and the inner VLAN ID 115 (dot1q). |
| (config-if)# rewrite pop-2tag | Configure rewrite with action pop-2tag. |
| (config-if)# access-if-vpls | Access VPLS under the sub-interface. |
| (config-acc-if-vpls)# mpls-vpls v108 | Associate the VPLS instance with the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE2: POP

| | |
|---|--|
| (config)#interface xe0.105 switchport | Create L2 sub-interface xe0.105 on the physical interface xe0. |
| (config-if)# encapsulation dot1ad 3001 inner-dot1q 2001 | Configure encapsulation with outer VLAN ID 3001 (dot1ad) and inner VLAN ID 2001 (dot1q). |
| (config-if)# rewrite pop | Configure rewrite with action pop. |
| (config-if)# access-if-vpls | Access VPLS under the sub-interface. |
| (config-acc-if-vpls)# mpls-vpls v105 | Associate the VPLS instance with the attachment circuit interface. |
| (config -acc-if-vpls)#exit | Exit from access mode. |

PE2: XLATE

| | |
|---|---|
| (config)# interface xe0.106 switchport | Create L2 sub-interface xe0.106 on the physical interface xe0. |
| (config-if)# encapsulation dot1ad 510 inner-dot1q 511 | Configure encapsulation with outer VLAN ID 510 (dot1ad) and inner VLAN ID 511(dot1q). |
| (config-if)# rewrite translate 2-to-1 dot1q 251 | Configure VLAN tag (251) translation on the interface. |
| (config-if)# access-if-vpls | Access VPLS under the sub-interface. |
| (config-acc-if-vpls)# mpls-vpls v106 | Associate the VPLS instance with the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE2: PUSH

| | |
|---|--|
| (config)#interface xe0.107 switchport | Create L2 sub-interface xe0.107 on the physical interface xe0. |
| (config-if)# encapsulation dot1ad 301 inner-dot1q 310 | Configure encapsulation with the outer VLAN ID 301 (dot1ad) and the inner VLAN ID 310 (dot1q). |
| (config-if)# rewrite push dot1q 321 | Configure the interface to push a single VLAN tag with VLAN ID 321. |
| (config-if)# access-if-vpls | Access VPLS under the sub-interface. |

| | |
|--------------------------------------|--|
| (config-acc-if-vpls)# mpls-vpls v107 | Associate the VPLS instance with the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

PE2: POP-2TAG

| | |
|---|--|
| (config)#interface xe0.108 switchport | Create L2 sub-interface xe0.108 on the physical interface xe0. |
| (config-if)# encapsulation dot1ad 105 inner-dot1q 115 | Configure encapsulation with the outer VLAN ID 105 (dot1ad) and the inner VLAN ID 115 (dot1q). |
| (config-if)# rewrite pop-2tag | Configure rewrite with action pop-2tag. |
| (config-if)# access-if-vpls | Access VPLS under the sub-interface. |
| (config-acc-if-vpls)# mpls-vpls v108 | Associate the VPLS instance with the attachment circuit interface. |
| (config-acc-if-vpls)#exit | Exit from access mode. |

Validation**PE1**

```

PE1#show bgp l2vpn vpls
VPLS-ID    VE-ID    Discovered-Peers  Route-Target
105         2        1                 100:105
106         2        1                 100:106
107         2        1                 100:107
108         2        1                 100:108

PE1#show bgp l2vpn vpls detail

VPLS ID: 105
VE-ID: 2
Discovered Peers: 1
Route-Target: 100:105
Local RD: 100:105
All Local Label Blocks:
[LB:25600, VBO:1, VBS:64]
Mesh Peers:
BGP Peer:2.2.2.2/32
VC Nbr Address:2.2.2.2, RD:100:105, VE-ID:1
VC Details: VC-ID:21
Local MTU:1500, Remote MTU:1500
Remote (LB:25600,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25600, Out Label:25601
PW Status:Established
VC Installed:Yes
VC Signaled Time: 00:21:31
Extended-Community Color:0

VPLS ID: 106
VE-ID: 2
Discovered Peers: 1
Route-Target: 100:106
Local RD: 100:106
All Local Label Blocks:
[LB:25664, VBO:1, VBS:64]
Mesh Peers:

```

```

BGP Peer:2.2.2.2/32
VC Nbr Address:2.2.2.2, RD:100:106, VE-ID:1
VC Details: VC-ID:21
Local MTU:1500, Remote MTU:1500
Remote (LB:25664,VBO:1,VBS:64) Local (LB:25664,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25664, Out Label:25665
PW Status:Established
VC Installed:Yes
VC Signaled Time: 00:21:15
Extended-Community Color:0

```

```

VPLS ID: 107
VE-ID: 2
Discovered Peers: 1
Route-Target: 100:107
Local RD: 100:107
All Local Label Blocks:
[LB:25728, VBO:1, VBS:64]
Mesh Peers:
BGP Peer:2.2.2.2/32
VC Nbr Address:2.2.2.2, RD:100:107, VE-ID:1
VC Details: VC-ID:21
Local MTU:1500, Remote MTU:1500
Remote (LB:25728,VBO:1,VBS:64) Local (LB:25728,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25728, Out Label:25729
PW Status:Established
VC Installed:Yes
VC Signaled Time: 00:21:01
Extended-Community Color:0

```

```

VPLS ID: 108
VE-ID: 2
Discovered Peers: 1
Route-Target: 100:108
Local RD: 100:108
All Local Label Blocks:
[LB:25792, VBO:1, VBS:64]
Mesh Peers:
BGP Peer:2.2.2.2/32
VC Nbr Address:2.2.2.2, RD:100:108, VE-ID:1
VC Details: VC-ID:21
Local MTU:1500, Remote MTU:1500
Remote (LB:25792,VBO:1,VBS:64) Local (LB:25792,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25792, Out Label:25793
PW Status:Established
VC Installed:Yes
VC Signaled Time: 00:20:31
Extended-Community Color:0

```

```

PE1#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| | | | | | | | | | |
|-----|---------|----------|---|-------|-------|-------|------|---|---|
| 105 | 2.2.2.2 | 25601 | | 25600 | ce2/1 | 25601 | 2/Up | 1 | B |
| GP | Active | 00:22:12 | - | | | | | | |
| 106 | 2.2.2.2 | 25601 | | 25664 | ce2/1 | 25665 | 2/Up | 2 | B |
| GP | Active | 00:21:56 | - | | | | | | |
| 107 | 2.2.2.2 | 25601 | | 25728 | ce2/1 | 25729 | 2/Up | 3 | B |
| GP | Active | 00:21:42 | - | | | | | | |
| 108 | 2.2.2.2 | 25601 | | 25792 | ce2/1 | 25793 | 2/Up | 4 | B |
| GP | Active | 00:21:12 | - | | | | | | |

```
PE1#show mpls vpls detail
```

```
Virtual Private LAN Service Instance: v105, ID: 105
SIG-Protocol: BGP
  Route-Distinguisher :100:105
  Route-Target :100:105
  VE-ID :2
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Raw
MAC Withdrawal:
```

```
Configured interfaces:
Interface: xe2.105
Status: Up
Subinterface Match Criteria(s) :
dot1ad 3001 inner-dot1q 2001
```

```
Mesh Peers:
  2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:25:14)
```

```
Virtual Private LAN Service Instance: v106, ID: 106
SIG-Protocol: BGP
  Route-Distinguisher :100:106
  Route-Target :100:106
  VE-ID :2
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Raw
MAC Withdrawal:
```

```
Configured interfaces:
Interface: xe2.106
Status: Up
Subinterface Match Criteria(s) :
dot1ad 510 inner-dot1q 511
```

```
Mesh Peers:
  2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:24:58)
```

```
Virtual Private LAN Service Instance: v107, ID: 107
SIG-Protocol: BGP
  Route-Distinguisher :100:107
  Route-Target :100:107
  VE-ID :2
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
```

```

Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe2.107
Status: Up
Subinterface Match Criteria(s) :
dot1ad 301 inner-dot1q 310

Mesh Peers:
2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:24:44)

Virtual Private LAN Service Instance: v108, ID: 108
SIG-Protocol: BGP
Route-Distinguisher :100:108
Route-Target :100:108
VE-ID :2
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe2.108
Status: Up
Subinterface Match Criteria(s) :
dot1ad 105 inner-dot1q 115

Mesh Peers:
2.2.2.2 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:24:14)

```

PE2

```

PE2#show bgp l2vpn vpls
VPLS-ID      VE-ID      Discovered-Peers  Route-Target
105          1          1                 100:105
106          1          1                 100:106
107          1          1                 100:107
108          1          1                 100:108

PE2#show bgp l2vpn vpls detail

VPLS ID: 105
VE-ID: 1
Discovered Peers: 1
Route-Target: 100:105
Local RD: 100:105
All Local Label Blocks:
[LB:25600, VBO:1, VBS:64]
Mesh Peers:
BGP Peer:1.1.1.1/32
VC Nbr Address:1.1.1.1, RD:100:105, VE-ID:2
VC Details: VC-ID:12
Local MTU:1500, Remote MTU:1500
Remote (LB:25600,VBO:1,VBS:64) Local (LB:25600,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25601, Out Label:25600
PW Status:Established

```

```
VC Installed:Yes
VC Signaled Time: 00:23:25
Extended-Community Color:0
```

```
VPLS ID: 106
VE-ID: 1
Discovered Peers: 1
Route-Target: 100:106
Local RD: 100:106
All Local Label Blocks:
[LB:25664, VBO:1, VBS:64]
Mesh Peers:
BGP Peer:1.1.1.1/32
VC Nbr Address:1.1.1.1, RD:100:106, VE-ID:2
VC Details: VC-ID:12
Local MTU:1500, Remote MTU:1500
Remote (LB:25664,VBO:1,VBS:64) Local (LB:25664,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25665, Out Label:25664
PW Status:Established
VC Installed:Yes
VC Signaled Time: 00:23:11
Extended-Community Color:0
```

```
VPLS ID: 107
VE-ID: 1
Discovered Peers: 1
Route-Target: 100:107
Local RD: 100:107
All Local Label Blocks:
[LB:25728, VBO:1, VBS:64]
Mesh Peers:
BGP Peer:1.1.1.1/32
VC Nbr Address:1.1.1.1, RD:100:107, VE-ID:2
VC Details: VC-ID:12
Local MTU:1500, Remote MTU:1500
Remote (LB:25728,VBO:1,VBS:64) Local (LB:25728,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25729, Out Label:25728
PW Status:Established
VC Installed:Yes
VC Signaled Time: 00:22:55
Extended-Community Color:0
```

```
VPLS ID: 108
VE-ID: 1
Discovered Peers: 1
Route-Target: 100:108
Local RD: 100:108
All Local Label Blocks:
[LB:25792, VBO:1, VBS:64]
Mesh Peers:
BGP Peer:1.1.1.1/32
VC Nbr Address:1.1.1.1, RD:100:108, VE-ID:2
VC Details: VC-ID:12
Local MTU:1500, Remote MTU:1500
Remote (LB:25792,VBO:1,VBS:64) Local (LB:25792,VBO:1,VBS:64)
LB sent on known VEID:Yes
In Label:25793, Out Label:25792
PW Status:Established
VC Installed:Yes
VC Signaled Time: 00:22:27
Extended-Community Color:0
```

```
PE2#show mpls vpls mesh
```

(m) - Service mapped over multipath transport

| | | | | | | | | |
|-----|---------|----------|-------|------|-------|------|----|---|
| 105 | 1.1.1.1 | 301 | 25601 | xe13 | 25600 | 2/Up | 6 | B |
| GP | Active | 00:24:04 | - | | | | | |
| 106 | 1.1.1.1 | 301 | 25665 | xe13 | 25664 | 2/Up | 8 | B |
| GP | Active | 00:23:50 | - | | | | | |
| 107 | 1.1.1.1 | 301 | 25729 | xe13 | 25728 | 2/Up | 9 | B |
| GP | Active | 00:23:34 | - | | | | | |
| 108 | 1.1.1.1 | 301 | 25793 | xe13 | 25792 | 2/Up | 10 | B |
| GP | Active | 00:23:06 | - | | | | | |

PE2#show mpls vpls detail

Virtual Private LAN Service Instance: v105, ID: 105

SIG-Protocol: BGP

Route-Distinguisher :100:105

Route-Target :100:105

VE-ID :1

Attachment-Circuit: UP

Learning: Enabled

Control-Word: Disabled

Flow Label Status: Disabled, Direction: None, Static: No

Group ID: 0, Configured MTU: 1500

Description: none

service-tpid: dot1.ad

Operating mode: Raw

MAC Withdrawal:

Configured interfaces:

Interface: xe0.105

Status: Up

Subinterface Match Criteria(s) :

dot1ad 3001 inner-dot1q 2001

Mesh Peers:

1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:24:28)

Virtual Private LAN Service Instance: v106, ID: 106

SIG-Protocol: BGP

Route-Distinguisher :100:106

Route-Target :100:106

VE-ID :1

Attachment-Circuit: UP

Learning: Enabled

Control-Word: Disabled

Flow Label Status: Disabled, Direction: None, Static: No

Group ID: 0, Configured MTU: 1500

Description: none

service-tpid: dot1.ad

Operating mode: Raw

MAC Withdrawal:

Configured interfaces:

Interface: xe0.106

Status: Up

Subinterface Match Criteria(s) :

dot1ad 510 inner-dot1q 511

Mesh Peers:

1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:24:14)

Virtual Private LAN Service Instance: v107, ID: 107

SIG-Protocol: BGP

Route-Distinguisher :100:107

Route-Target :100:107

VE-ID :1

```
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe0.107
Status: Up
Subinterface Match Criteria(s) :
dotlad 301 inner-dot1q 310

Mesh Peers:
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:23:58)

Virtual Private LAN Service Instance: v108, ID: 108
SIG-Protocol: BGP
Route-Distinguisher :100:108
Route-Target :100:108
VE-ID :1
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1.ad
Operating mode: Raw
MAC Withdrawal:

Configured interfaces:
Interface: xe0.108
Status: Up
Subinterface Match Criteria(s) :
dotlad 105 inner-dot1q 115

Mesh Peers:
1.1.1.1 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:23:30)
```

Hierarchical VPLS

Overview

A Virtual Private LAN Service (VPLS) enables multipoint to multipoint communication, creating LAN-like connectivity between customers' sites. However, the typical full mesh topology required for LAN emulation can be impractical in large networks. To address this, Hierarchical VPLS (H-VPLS) introduces a hierarchical approach using a spoke-PW (pseudowire) type. Unlike the standard mesh-PW, the spoke-PW facilitates traffic between hierarchical levels, offering a more scalable solution for VPLS networks.

H-VPLS Redundancy Characteristics

In a Virtual Private LAN Service (VPLS) network, when a node connects through a spoke-PW, a single point of failure arises. In the event of a connection failure to the VPLS mesh or a failure within the PE-rs node, the spoke device experiences a complete loss of connectivity. To address this, PW redundancy is implemented, configuring a secondary path that activates if the primary path fails. The MTU-s is configured with a primary spoke-PW connected to PE1-rs and a secondary spoke-PW connected to PE2-rs. During normal operation, the primary spoke-PW is active, but in case of failure, the MTU-s can switch to the standby spoke-PW for continued connectivity, aiming for sub-second convergence times with potential MAC flush-related traffic loss.

Benefits

Hierarchical VPLS (H-VPLS) is introduced to address scalability challenges associated with the traditional VPLS (Virtual Private LAN Service) architecture. It introduces a hierarchical approach that enhances scalability, reduces configuration complexity, optimizes traffic flow, and improves overall network efficiency and fault tolerance.

Limitations

- Automatic revertive cases from secondary to primary will not be supported.
- MAC Address Withdrawal feature will not be supported in release 6.5.2.
- Convergence on redundancy may require bidirectional traffic or MAC aging.

Prerequisites

- The `block-mesh-spoke-on-all-ac-down` and `ignore-ac-spoke-state` commands are optional and mutually exclusive, meaning only one can be applied at a time, or neither. By default, neither command is applied. If one of commands is applied, applying the other will make it the active one. To remove a command, use the `no` prefix.

```
signaling ldp
(block-mesh-spoke-on-all-ac-down | ignore-ac-spoke-state)
(no block-mesh-spoke-on-all-ac-down | no ignore-ac-spoke-state)
```

- **Define Interfaces and Loopback Addresses:**
Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

```
!
```

```
interface lo
  ip address 127.0.0.1/8
  ip address 2.2.2.2/32 secondary
  ipv6 address ::1/128

interface xe14
  ip address 30.1.1.2/24
```

- **Configure IGP for Dynamic Routing:** Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

ISIS Configuration

```
router isis 1
  is-type level-2-only
  metric-style wide
  microloop-avoidance level-2
  mpls traffic-eng router-id 2.2.2.2
  mpls traffic-eng level-2
  capability cspf
  dynamic-hostname
  bfd all-interfaces
  net 49.0000.0000.0002.00
  passive-interface lo
!
interface xe14
  mpls ldp-igp sync isis level-2
  isis network point-to-point
  ip router isis 1
```

OSPF Configuration


```
router ospf 1
  ospf router-id 2.2.2.2
  network 2.2.2.2/32 area 0.0.0.0
  network 30.1.1.0/24 area 0.0.0.0!
!
interface xe14
  ip ospf network point-to-point
```

Configuration for H-VPLS with Redundancy

Configure various nodes within the topology to set up a H-VPLS session.

Topology

This sample topology provides basic connectivity and routing between the devices.

Figure 80. H-VPLS Configuration with Redundancy

Configure H-VPLS on PE1 Router

Follow the steps to configure the H-VPLS on PE1 router:

1. Configure router LDP.

```
PE1(config)#router ldp
PE1(config-router)# router-id 2.2.2.2
PE1(config-router)# transport-address ipv4 2.2.2.2
```

2. Configure targeted-peer under router LDP.

```
PE1(config-router)# targeted-peer ipv4 5.5.5.5
PE1(config-router-targeted-peer)# exit-targeted-peer-mode
PE1(config-router)# targeted-peer ipv4 3.3.3.3
PE1(config-router-targeted-peer)# exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
PE1(config)#interface xe14
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching

PE1(config)#interface xe26
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE1(config)#mpls vpls vpls2000 2000
PE1(config-vpls)# signaling ldp
PE1(config-vpls-sig)# vpls-peer 3.3.3.3
PE1(config-vpls-sig)# vpls-peer 5.5.5.5
PE1(config-vpls-sig)# exit-signaling
PE1(config-vpls)# exit-vpls
PE1(config)#
```

5. Configure sub-interface and attach vpls-instance to sub-interface.

```
PE1(config)#
PE1(config)#interface xe16.2000 switchport
PE1(config-if)# encapsulation dot1q 2000
PE1(config-if)# access-if-vpls
PE1(config-acc-if-vpls)# mpls-vpls vpls2000
PE1(config-acc-if-vpls)#
```

Configure H-VPLS on PE2 (Primary Hub)

Follow the steps to configure the H-VPLS on PE2 (Primary Hub):

1. Configure router LDP.

```
PE2(config)#router ldp
PE2(config-router)# router-id 5.5.5.5
PE2(config-router)# transport-address ipv4 5.5.5.5
```

2. Configure targeted-peer under router LDP.

```
PE2(config)#router ldp
PE2(config-router)# targeted-peer ipv4 2.2.2.2
PE2(config-router-targeted-peer)# exit-targeted-peer-mode
PE2(config-router)# targeted-peer ipv4 3.3.3.3
PE2(config-router-targeted-peer)# exit-targeted-peer-mode
PE2(config-router)#
```

3. Enable LDP and label-switching for core interface.

```
PE2(config)#interface xe1
PE2(config-if)# enable-ldp ipv4
PE2(config-if)#label-switching

PE2(config)#interface xe12
PE2(config-if)# enable-ldp ipv4
PE2(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE2(config)#mpls vpls vpls2000 2000
PE2(config-vpls)# signaling ldp
PE2(config-vpls-sig)# vpls-peer 2.2.2.2
PE2(config-vpls-sig)# vpls-peer 3.3.3.3
PE2(config-vpls-sig)# exit-signaling
PE2(config-vpls)# exit-vpls
PE2(config)#
```

5. Configure L2-ckt.

```
PE2(config)#mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
PE2(config-pseudowire)#
```

6. Attach L2-ckt under vpls instance.

```
PE2(config)#mpls vpls vpls2000 2000
PE2(config-vpls)#vpls-vc vc2000
PE2(config-vpls-spoke)#
```

Configure H-VPLS on PE3 (Secondary Hub)

Follow the steps to configure the H-VPLS on PE3 (Secondary Hub):

1. Configure router LDP.

```
PE3(config)#router ldp
PE3(config-router)# router-id 3.3.3.3
PE3(config-router)# transport-address ipv4 3.3.3.3
```

2. Configure targeted-peer under router LDP.

```
PE3(config)#router ldp
PE3(config-router)# targeted-peer ipv4 2.2.2.2
PE3(config-router-targeted-peer)# exit-targeted-peer-mode
PE3(config-router)# targeted-peer ipv4 5.5.5.5
PE3(config-router-targeted-peer)# exit-targeted-peer-mode
PE3(config-router)#
```

3. Enable LDP and label-switching for core interface.

```
PE3(config)#interface xe23
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching

PE3(config)#interface xe26
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE3(config)#mpls vpls vpls2000 2000
PE3(config-vpls)# signaling ldp
PE3(config-vpls-sig)# vpls-peer 2.2.2.2
PE3(config-vpls-sig)# vpls-peer 5.5.5.5
PE3(config-vpls-sig)# exit-signaling
PE3(config-vpls)# exit-vpls
PE3(config)#
```

5. Configure L2-ckt.

```
PE3(config)#mpls l2-circuit vc2001 2223 8.8.8.8 mode raw
PE3(config-pseudowire)#
```

6. Attach L2-ckt under vpls instance.

```
PE3 (config)#mpls vpls vpls2000 2000
PE3(config-vpls)#vpls-vc vc2001
PE3(config-vpls-spoke)#
```

Configure H-VPLS on Spoke Router

Follow the steps to configure the H-VPLS on Spoke router:

1. Configure router LDP.

```
Spoke(config)#router ldp
Spoke(config-router)# router-id 8.8.8.8
Spoke(config-router)# transport-address ipv4 8.8.8.8
```

2. Configure targeted-peer under router LDP.

```
Spoke(config-router)# targeted-peer ipv4 5.5.5.5
Spoke(config-router-targeted-peer)# exit-targeted-peer-mode
Spoke(config-router)# targeted-peer ipv4 3.3.3.3
Spoke(config-router-targeted-peer)# exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
Spoke(config)#interface xe12
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching

Spoke(config)#interface xe25
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching
```

4. Configure VPLS instance.

```
Spoke (config)#mpls vpls vpls2000 2000
Spoke (config-vpls)#
```

5. Configure L2-ckt.

```
Spoke(config)#mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
Spoke(config-pseudowire)#!
Spoke(config-pseudowire)#mpls l2-circuit vc2001 2223 3.3.3.3 mode raw
Spoke(config-pseudowire)#
```

6. Configure Primary and secondary spoke under vpls instance.

```
Spoke(config)#mpls vpls vpls2000 2000
Spoke(config-vpls)#vpls-vc vc2000
Spoke(config-vpls-spoke)# secondary vc2001
Spoke(config-vpls-spoke)# exit-spoke
Spoke(config-vpls)# exit-vpls
Spoke(config)#
```

7. Configure sub-interface and attach vpls-instance to sub-interface.

```
Spoke(config)#
Spoke(config)#interface xe26.2000 switchport
Spoke(config-if)# encapsulation dot1q 2000
Spoke(config-if)# access-if-vpls
Spoke(config-acc-if-vpls)# mpls-vpls vpls2000
Spoke(config-acc-if-vpls)#
```

Running Configuration**PE1**

```
router ldp
router-id 2.2.2.2
targeted-peer ipv4 3.3.3.3
exit-targeted-peer-mode
targeted-peer ipv4 5.5.5.5
transport-address ipv4 2.2.2.2
!
interface xe14
enable-ldp ipv4
!
interface xe26
enable-ldp ipv4
!
mpls vpls vpls2000 2000
signaling ldp
vpls-peer 3.3.3.3
vpls-peer 5.5.5.5
exit-signaling
exit-vpls
!
interface xe16.2000 switchport
access-if-vpls
mpls-vpls vpls2000
```

PE2

```
router ldp
targeted-peer ipv4 2.2.2.2
exit-targeted-peer-mode
targeted-peer ipv4 3.3.3.3
exit-targeted-peer-mode
transport-address ipv4 5.5.5.5
!
mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
!
mpls vpls vpls2000 2000
vpls-vc vc2000
exit-spoke
signaling ldp
vpls-peer 2.2.2.2
vpls-peer 3.3.3.3
exit-signaling
exit-vpls
```

PE3

```

router ldp
  targeted-peer ipv4 2.2.2.2
  exit-targeted-peer-mode
  targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
  transport-address ipv4 3.3.3.3
!
mpls l2-circuit vc2001 2223 8.8.8.8 mode raw
!
mpls vpls vpls2000 2000
  vpls-vc vc2001
  exit-spoke
  signaling ldp
  vpls-peer 2.2.2.2
  vpls-peer 5.5.5.5
  exit-signaling
  exit-vpls

```

Spoke Router

```

router ldp
  router-id 8.8.8.8
  targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
  targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
  transport-address ipv4 8.8.8.8
!
mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
!
mpls l2-circuit vc2001 2223 3.3.3.3 mode raw
!
mpls vpls vpls2000 2000
  vpls-vc vc2000
  secondary vc2001
  exit-spoke
  exit-vpls
!
interface xe26.2000 switchport
  access-if-vpls
  mpls-vpls vpls2000

```

Validation

Validate the show output after configuration as shown below.
Verify vpls mesh are up between PE1 and Hub Nodes

```

PE1#sho mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX | |
|--------------|-----------|--------------|----------|--------------|-----------|---------|----------|---|
| SIG-Protocol | Status | UpTime | | | | | | |
| 2000 | 3.3.3.3 | 29447 | 28164 | xe26 | 27532 | 2/Up | 3 | L |
| DP | Active | 2d12h08m | | | | | | |
| 2000 | 5.5.5.5 | 31364 | 28162 | xe14 | 26883 | 2/Up | 4 | |
| LDP | Active | 2d12h04m | | | | | | |

```

PE2#sho mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|---------|-----------|--------------|----------|--------------|-----------|---------|----------|
|---------|-----------|--------------|----------|--------------|-----------|---------|----------|

| SIG-Protocol | Status | UpTime | | | | | | | |
|--------------|---------|----------|-------|-----|-------|------|---|--|-----|
| 2000 | 2.2.2.2 | 29446 | 26883 | xe1 | 28162 | 2/Up | 3 | | LDP |
| | Active | 2d12h05m | | | | | | | |
| 2000 | 3.3.3.3 | 31367 | 26884 | xe1 | 27528 | 2/Up | 4 | | L |
| DP | Active | 2d12h15m | | | | | | | |

```

PE3#sho mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | | | | | |
| 2000 | 2.2.2.2 | 29440 | 27532 | xe26 | 28164 | 2/Up | 3 |
| LDP | Active | 2d12h10m | | | | | |
| 2000 | 5.5.5.5 | 31363 | 27528 | xe26 | 26884 | 2/Up | 4 |
| LDP | Active | 2d12h16m | | | | | |

Verify vpls spoke between Hub and Spoke

```

PE2#sho mpls vpls spoke
VPLS-ID   Virtual Circuit  Tunnel-Label In-Label   Network-Intf Out-Label  Lkps/St   Secondary
2000      vc2000              29443        26882      xe1          26886     2/Up      ---

```

```

PE3#show mpls vpls spoke
VPLS-ID   Virtual Circuit  Tunnel-Label In-Label   Network-Intf Out-Label  Lkps/St   Secondary
2000      vc2001          N/A          27527      N/A          26883     0/Dn      ---

```

```

Spoke#show mpls vpls spoke
VPLS-ID   Virtual Circuit  Tunnel-Label In-Label   Network-Intf Out-Label  Lkps/St   Secondary
2000      vc2000          29440        26886      xe12         26882     2/Up      vc2001
2000      vc2001          N/A          26883      N/A          27527     0/Dn      ---

```

Verify H-vpls session on Hub and spoke:

```

PE2#show mpls vpls vpls2000
Virtual Private LAN Service Instance: vpls2000, ID: 2000
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Ignoring AC interface and spoke-VC state

```

Configured interfaces:
None

```

Mesh Peers:
  2.2.2.2 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h13m)
  3.3.3.3 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h22m)
Spoke Peers:
  vc2000 (Up) (UpTime 01:31:27)

```

```

PE3#show mpls vpls vpls2000
Virtual Private LAN Service Instance: vpls2000, ID: 2000
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q

```



```

Operating mode: Raw
Ignoring AC interface and spoke-VC state

Configured interfaces:
  None

Mesh Peers:
  2.2.2.2 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h16m)
  5.5.5.5 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h22m)
Spoke Peers:
  vc2001 (Dn) (Reason: VC on standby)

Spoke#show mpls vpls vpls2000
Virtual Private LAN Service Instance: vpls2000, ID: 2000
SIG-Protocol: N/A
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw

Configured interfaces:
  Interface: xe26.2000
  Status: Up
  Subinterface Match Criteria(s) :
  dot1q 2000

Spoke Peers:
  vc2000 (Up) (UpTime 01:31:33)
  Secondary: vc2001 (Dn) (Reason: VC on standby)

```


Configuration for H-VPLS without Redundancy

Configure various nodes within the topology to set up a H-VPLS session.

Topology

This sample topology provides basic connectivity and routing between the devices.

Figure 81. H-VPLS Configuration without Redundancy



Configure H-VPLS on PE1 Router

Follow the steps to configure the H-VPLS on PE1 router:

1. Configure router LDP.

```

PE1(config)#router ldp
PE1(config-router)# router-id 2.2.2.2
PE1(config-router)# transport-address ipv4 2.2.2.2

```

2. Configure targeted-peer under router LDP.

```
PE1(config-router)# targeted-peer ipv4 5.5.5.5
PE1(config-router-targeted-peer)# exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
PE1(config)#interface xe14
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE1(config)#mpls vpls vpls2000 2000
PE1(config-vpls)# signaling ldp
PE1(config-vpls-sig)# vpls-peer 5.5.5.5
PE1(config-vpls-sig)# exit-signaling
PE1(config-vpls)# exit-vpls
PE1(config)#
```

5. Configure sub-interface and attach vpls-instance to sub-interface

```
PE1(config)#
PE1(config)#interface xe16.2000 switchport
PE1(config-if)# encapsulation dot1q 2000
PE1(config-if)# access-if-vpls
PE1(config-acc-if-vpls)# mpls-vpls vpls2000
PE1(config-acc-if-vpls)#
```

Configure H-VPLS on Hub Router

Follow the steps to configure the H-VPLS on Hub router:

1. Configure router LDP.

```
Hub(config)#router ldp
Hub(config-router)# router-id 5.5.5.5
Hub(config-router)# transport-address ipv4 5.5.5.5
```

2. Configure targeted-peer under router LDP.

```
Hub(config-router)# targeted-peer ipv4 2.2.2.2
Hub(config-router-targeted-peer)# exit-targeted-peer-mode
R5-P5(config-router)# targeted-peer ipv4 8.8.8.8
R5-P5(config-router-targeted-peer)#
```

3. Enable LDP and label-switching for core interface.

```
Hub(config)#interface xe1
Hub(config-if)# enable-ldp ipv4
Hub(config-if)#label-switching

Hub(config)#interface xe12
Hub(config-if)# enable-ldp ipv4
Hub(config-if)#label-switching
```

4. Configure VPLS instance.

```
Hub(config)#mpls vpls vpls2000 2000
Hub(config-vpls)# signaling ldp
Hub(config-vpls-sig)# vpls-peer 2.2.2.2
Hub(config-vpls-sig)# exit-signaling
Hub(config-vpls)# exit-vpls
Hub(config)#
```

5. Configure L2-ckt.

```
Hub (config)#mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
Hub (config-pseudowire)#
```

6. Attach L2-ckt under vpls instance.

```

Hub (config)#mpls vpls vpls2000 2000
Hub (config-vpls)#vpls-vc vc2000
Hub (config-vpls-spoke)#

```

Configure H-VPLS on Spoke Router

Follow the steps to configure the H-VPLS on Spoke router:

1. Configure router LDP.

```

Spoke (config)#router ldp
Spoke (config-router)# router-id 8.8.8.8
Spoke (config-router)# transport-address ipv4 8.8.8.8

```

2. Configure targeted-peer under router LDP.

```

Spoke (config-router)# targeted-peer ipv4 5.5.5.5
Spoke (config-router-targeted-peer)# exit-targeted-peer-mode

```

3. Enable LDP and label-switching for core interface.

```

Spoke (config)#interface xe12
Spoke (config-if)# enable-ldp ipv4
Spoke (config-if)#label-switching

```

4. Configure VPLS instance.

```

Spoke (config)#mpls vpls vpls2000 2000
Spoke (config-vpls)#

```

5. Configure L2-ckt.

```

Spoke (config)#mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
Spoke (config-pseudowire)#

```

6. Attach L2-ckt under vpls instance.

```

Spoke (config)#mpls vpls vpls2000 2000
Spoke (config-vpls)#vpls-vc vc2000
Spoke (config-vpls-spoke)#

```

7. Configure sub-interface and attach vpls-instance to sub-interface.

```

Spoke (config)#
Spoke (config)#interface xe26.2000 switchport
Spoke (config-if)# encapsulation dot1q 2000
Spoke (config-if)# access-if-vpls
Spoke (config-acc-if-vpls)# mpls-vpls vpls2000
Spoke (config-acc-if-vpls)#

```

Running Configuration

PE1

```

router ldp
router-id 2.2.2.2
targeted-peer ipv4 5.5.5.5
exit-targeted-peer-mode
transport-address ipv4 2.2.2.2
!
interface xe14
enable-ldp ipv4
!
mpls vpls vpls2000 2000
signaling ldp
vpls-peer 5.5.5.5

```

```

    exit-signaling
  exit-vpls
!
interface xe16.2000 switchport
  access-if-vpls
    mpls-vpls vpls2000

```

Hub

```

router ldp
  targeted-peer ipv4 2.2.2.2
  exit-targeted-peer-mode
    targeted-peer ipv4 8.8.8.8
    exit-targeted-peer-mode
  !
!
mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
!
mpls vpls vpls2000 2000
  vpls-vc vc2000
  exit-spoke
  signaling ldp
    vpls-peer 2.2.2.2
  exit-signaling
exit-vpls

```

Spoke

```

router ldp
  router-id 8.8.8.8
  targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
    transport-address ipv4 8.8.8.8
  !
!
mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
!
mpls vpls vpls2000 2000
  vpls-vc vc2000
  exit-spoke
exit-vpls
!
interface xe26.2000 switchport
  access-if-vpls
    mpls-vpls vpls2000

```

Validation

Validate the show output after configuration as shown below.
Verify vpls mesh are up between PE and Hub

```

PE1#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID INDEX | Peer Addr SIG-Protocol | Status | Tunnel-Label UpTime | In-Label | Network-Intf | Out-Label | Lkps/St | PW- |
|------------------|---------------------------|--------|------------------------|----------|--------------|-----------|---------|-----|
| 2000 | 5.5.5.5 | Active | 31364 2d10h36m | 28162 | xe14 | 26883 | 2/Up | 4 |

```

Hub#sho mpls vpls mesh
(m) - Service mapped over multipath transport

```

(e) - Service mapped over LDP ECMP

| VPLS-ID INDEX | Peer Addr SIG-Protocol | Status | Tunnel-Label UpTime | In-Label | Network-Intf | Out-Label | Lkps/St | PW- |
|------------------|---------------------------|--------|------------------------|----------|--------------|-----------|---------|-----|
| 2000 | 2.2.2.2 | Active | 29446 2d10h39m | 26883 | xe1 | 28162 | 2/Up | 3 |

Verify vpls spoke are up between Hub and Spoke

Hub#sho ldp mpls-l2-circuit

| Transport | Client | VC | VC | Local | Remote | Destination | Lo- |
|-----------|-----------|-----------|----------|----------|----------|-------------|------------|
| cal | | | Remote | | | | |
| VC ID | Binding | State | Type | VC Label | VC Label | Address | PW |
| Status | | PW Status | | | | | |
| 2222 | VPLS:2000 | UP | Ethernet | 26882 | 26886 | 8.8.8.8 | Forwarding |

Hub#sho mpls vpls spoke

| VPLS-ID | Virtual Circuit | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | Secondary |
|---------|-----------------|--------------|----------|--------------|-----------|---------|-----------|
| 2000 | vc2000 | 29443 | 26882 | ce4 | 26886 | 2/Up | --- |

Spoke#show ldp mpls-l2-circuit

| Transport | Client | VC | VC | Local | Remote | Destination | Lo- |
|-----------|-----------|-----------|----------|----------|----------|-------------|------------|
| cal | | | Remote | | | | |
| VC ID | Binding | State | Type | VC Label | VC Label | Address | PW |
| Status | | PW Status | | | | | |
| 2222 | VPLS:2000 | UP | Ethernet | 26886 | 26882 | 5.5.5.5 | Forwarding |

Spoke#show mpls vpls spoke

| VPLS-ID | Virtual Circuit | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | Secondary |
|---------|-----------------|--------------|----------|--------------|-----------|---------|-----------|
| 2000 | vc2000 | 29440 | 26886 | ce4 | 26882 | 2/Up | --- |

Verify H-vpls session on Hub and spoke:

Hub#show mpls vpls vpls2000

Virtual Private LAN Service Instance: vpls2000, ID: 2000
 SIG-Protocol: LDP
 Attachment-Circuit: UP
 Learning: Enabled
 Control-Word: Disabled
 Flow Label Status: Disabled, Direction: None, Static: No
 Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
 Description: none
 service-tpid: dot1.q
 Operating mode: Raw
 Ignoring AC interface and spoke-VC state

Configured interfaces:

None

Mesh Peers:

2.2.2.2 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d10h47m)
 3.3.3.3 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d10h56m)

Spoke Peers:

vc2000 (Up) (UpTime 00:05:48)

Spoke#show mpls vpls vpls2000

Virtual Private LAN Service Instance: vpls2000, ID: 2000
 SIG-Protocol: N/A
 Attachment-Circuit: UP
 Learning: Enabled
 Control-Word: Disabled
 Flow Label Status: Disabled, Direction: None, Static: No
 Group ID: 0, Configured MTU: 1500
 Description: none

```
service-tpid: dot1.q
Operating mode: Raw

Configured interfaces:
Interface: xe26.2000
Status: Up
Subinterface Match Criteria(s) :
dot1q 2000

Spoke Peers:
vc2000 (Up) (UpTime 00:07:47)
```

Commands for H-VPLS Configuration

The H-VPLS uses the following configuration commands.

vpls-vc

Use this command to add a spoke virtual circuit to VPLS domain hierarchically.

Use `no` parameter of this command to remove this configuration.

Command Syntax

```
vpls-vc NAME
(secondary NAME|)
(ethernet|vlan|)
```

Parameters

NAME

Specifies the name of the VPLS. It is a string that identifies the MPLS VC to add to the VPLS domain.

secondary

Specifies the name of the secondary spoke.

NAME

Specifies the name for the secondary spoke.

ethernet

Specifies the spoke type. Defaults to ethernet.

vlan

Specifies the spoke type.

Default

Disabled

Command Mode

VPLS mode

Applicability

Introduced before OcNOS version 1.3.

Modified the command prompt into a hierarchical structure from single line in the OcNOS version 6.5.1.

Example

Example for adding a spoke virtual circuit with VPLS name vc1 and secondary spoke vc2:

```
#configure terminal
(config)#mpls vpls vpls1 3000
(config-vpls)#vpls-vc vc1
(config-vpls-spoke)#secondary vc2
(config-vpls-spoke)#type ethernet
(config-vpls-spoke)#exit-spoke
(config-vpls)#exit
```

Example to remove the configuration of the spoke virtual circuit with VPLS name vc1:

```
#configure terminal
(config)#mpls vpls vpls1 3000
(config-vpls)#no vpls-vc vc1
(config-vpls)#exit
```

signaling

Use this command to set all mesh and spoke pseudowires to down when all access interfaces are down.

Use `ignore-ac-spoke-state` parameter of this command to remove this configuration.

Command Syntax

```
signaling ldp block-mesh-spoke-on-all-ac-down
```

```
signaling ignore-ac-spoke-state
```

Parameters

block-mesh-spoke-on-all-ac-down

(Optional) Controls the behavior of pseudowires (PWs) in a VPLS instance when all access interfaces associated with the VPLS instance are down.

ignore-ac-spoke-state

Ignores access interfaces and spoke pseudowires state and keep mesh pseudowires up.

Default

disabled

Command Mode

VPLS mode

Applicability

Introduced before OcNOS version 1.3.

Modified the command prompt into a hierarchical structure from single line in the OcNOS version 6.5.1.

Example

Example for setting up all mesh and spoke pseudowires to down when all access interfaces are down:

```
#configure terminal
(config)# mpls vpls test 100
(config-vpls)#signaling ldp
(config-vpls-sig)#block-mesh-spoke-on-all-ac-down
```

```
(config-vpls-sig)#exit
```

Example for setting up all mesh and spoke pseudowires to up:

```
#configure terminal
(config)# mpls vpls test 100
(config-vpls)#signaling ldp
(config-vpls-sig)#ignore-ac-spoke-state
(config-vpls-sig)#exit
```

MAC Withdrawal - VPLS/H-VPLS

Overview

The MAC withdrawal mechanism allows for faster convergence by immediately clearing dynamically learned MAC addresses upon detecting a switchover event, rather than waiting for the usual aging process. This process is triggered by specific network events or state changes, such as:

- Spoke-PW switchover: A path between two VPLS network elements, typically between MTU-s (Multi-Tenant Units) and Provider Edge (PE) routers, transitions from standby to active or vice versa.
- MLAG switchover: When a link failure or node failure occurs in a MLAG setup, leading to the activation of a previously standby link.
- VPWS Switchover: With VPWS when secondary becomes primary VC or primary becomes active with revertive mode, the mac withdrawal message would be originated from MTU-s.

The main goal is to ensure that MAC addresses learned from any given path (spoke or mesh) are promptly removed when a switchover occurs, preventing stale entries from affecting the network.

The MAC withdrawal message, defined as the LDP Positive MAC Flush message, is used to remove MAC address tables. This message can be triggered either by MTU-s or PE-rs depending on the network topology and the specific trigger event (such as a switchover or link failure).

Characteristics of MAC Withdrawal - VPLS/H-VPLS

- MAC Withdrawal is used to speed up network convergence during topology changes.
- It ensures that dynamically learned MAC addresses are flushed proactively instead of aging out, reducing traffic disruption.
- Triggers for MAC Withdrawal:
 - Spoke PW Switchover: When a Spoke Pseudowire (PW) transitions between active and standby.
 - MLAG Switchover: When a multi-chassis link aggregation (MLAG) event causes a switchover between active and standby devices.
 - Node Failure: When a PE-rs (Provider Edge router) or MTU-s (Multi-Tenant Unit switch) fails or reboots.
 - PW Status TLV: The Preferential Forwarding Status bit in an LDP PW Status TLV can signal a switchover.
 - Manual Intervention: Administrators can trigger MAC withdrawal manually for network maintenance.
- With VPWS when secondary becomes primary VC mac withdrawal message would be originated from MTU-s. Since revertive mode is supported the original active coming backup would also generate MAC withdrawal message towards the PE-rs.

Benefits

- Proactively removes outdated MAC entries, reducing the time taken for the network to stabilize.
- Ensures quick adaptation to topology changes without requiring manual intervention.
- Reduces packet flooding by clearing invalid MAC table entries.
- Prevents stale forwarding information, ensuring accurate traffic routing.
- Ensures seamless transition when a standby link becomes active.

- Enables compatibility with industry-standard VPLS/H-VPLS deployments.
- Prevents unnecessary flushing across Mesh PWs, maintaining stability.

Prerequisites

- Define Interfaces and Loopback Addresses:

Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 2.2.2.2/32 secondary  
  ipv6 address ::1/128  
  
interface xe14  
  ip address 30.1.1.2/24
```

- Configure IGP for Dynamic Routing:

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

- ISIS Configuration:

```
router isis 1  
  is-type level-2-only  
  metric-style wide  
  microloop-avoidance level-2  
  mpls traffic-eng router-id 2.2.2.2  
  mpls traffic-eng level-2  
  capability cspf  
  dynamic-hostname  
  bfd all-interfaces  
  net 49.0000.0000.0002.00  
  passive-interface lo  
  !  
interface xe14  
  mpls ldp-igp sync isis level-2  
  isis network point-to-point  
  ip router isis 1
```

- OSPF Configuration:


```
router ospf 1  
  ospf router-id 2.2.2.2  
  network 2.2.2.2/32 area 0.0.0.0  
  network 30.1.1.0/24 area 0.0.0.0!  
  !  
interface xe14  
  ip ospf network point-to-point
```

Configuration

Topology

This topology represents MAC withdrawal for VPLS/H-VPLS with Primary and Secondary Hub architecture.

Figure 82. MAC Withdrawal for VPLS/H-VPLS Topology



PE1

1. Configure router LDP.

```

PE1(config)#router ldp
PE1(config-router)# router-id 2.2.2.2
PE1(config-router)# transport-address ipv4 2.2.2.2
  
```

2. Configure targeted-peer under router LDP.

```

PE1(config-router)# targeted-peer ipv4 5.5.5.5
PE1(config-router-targeted-peer)# exit-targeted-peer-mode
PE1(config-router)# targeted-peer ipv4 3.3.3.3
PE1(config-router-targeted-peer)# exit-targeted-peer-mode
  
```

3. Enable LDP and label-switching for core interface.

```

PE1(config)#interface xe14
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching

PE1(config)#interface xe26
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching
  
```

4. Configure VPLS instance.

```

PE1(config)#mpls vpls vpls2000 2000
PE1(config-vpls)# signaling ldp
PE1(config-vpls-sig)# vpls-peer 3.3.3.3
PE1(config-vpls-sig)# vpls-peer 5.5.5.5
PE1(config-vpls-sig)# exit-signaling
PE1(config-vpls)# exit-vpls
PE1(config)#
  
```

5. Configure sub-interface and attach vpls-instance to sub-interface.

```
PE1(config)#
PE1(config)#interface xe16.2000 switchport
PE1(config-if)# encapsulation dot1q 2000
PE1(config-if)# access-if-vpls
PE1(config-acc-if-vpls)# mpls-vpls vpls2000
PE1(config-acc-if-vpls)#
```

PE2 (Primary Hub)

1. Configure router LDP.

```
PE2(config)#router ldp
PE2(config-router)# router-id 5.5.5.5
PE2(config-router)# transport-address ipv4 5.5.5.5
```

2. Configure targeted-peer under router LDP.

```
PE2(config)#router ldp
PE2(config-router)# targeted-peer ipv4 2.2.2.2
PE2(config-router-targeted-peer)# exit-targeted-peer-mode
PE2(config-router)# targeted-peer ipv4 3.3.3.3
PE2(config-router-targeted-peer)# exit-targeted-peer-mode
PE2(config-router)#
```

3. Enable LDP and label-switching for core interface

```
PE2(config)#interface xe1
PE2(config-if)# enable-ldp ipv4
PE2(config-if)#label-switching

PE2(config)#interface xe12
PE2(config-if)# enable-ldp ipv4
PE2(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE2(config)#mpls vpls vpls2000 2000
PE2(config-vpls)# signaling ldp
PE2(config-vpls-sig)# vpls-peer 2.2.2.2
PE2(config-vpls-sig)# vpls-peer 3.3.3.3
PE2(config-vpls-sig)# exit-signaling
PE2(config-vpls)# exit-vpls
PE2(config)#
```

5. Configure L2-ckt.

```
PE2 (config)#mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
PE2(config-pseudowire)#
```

6. Attach L2-ckt under vpls instance.

```
PE2 (config)#mpls vpls vpls2000 2000
PE2(config-vpls)#vpls-vc vc2000
PE2(config-vpls-spoke)#
```

7. Configure mac-withdrawal under VPLS instance.

```
PE2(config-vpls-spoke)#mpls vpls vpls2000 2000
PE2(config-vpls)# vpls-vc vc2000
PE2(config-vpls-spoke)# exit-spoke
PE2(config-vpls)# mac-withdrawal flush-propagate-spoke-to-mesh
PE2(config-vpls)#exit
PE2(config)#
```

PE3 (Secondary Hub)

1. Configure router LDP.

```
PE3(config)#router ldp
PE3(config-router)# router-id 3.3.3.3
PE3(config-router)# transport-address ipv4 3.3.3.3
```

2. Configure targeted-peer under router LDP.

```
PE3(config)#router ldp
PE3(config-router)# targeted-peer ipv4 2.2.2.2
PE3(config-router-targeted-peer)# exit-targeted-peer-mode
PE3(config-router)# targeted-peer ipv4 5.5.5.5
PE3(config-router-targeted-peer)# exit-targeted-peer-mode
PE3(config-router)#
```

3. Enable LDP and label-switching for core interface.

```
PE3(config)#interface xe23
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching
```

```
PE3(config)#interface xe26
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE3(config)#mpls vpls vpls2000 2000
PE3(config-vpls)# signaling ldp
PE3(config-vpls-sig)# vpls-peer 2.2.2.2
PE3(config-vpls-sig)# vpls-peer 5.5.5.5
PE3(config-vpls-sig)# exit-signaling
PE3(config-vpls)# exit-vpls
PE3(config)#
```

5. Configure L2-ckt.

```
PE3(config)#mpls l2-circuit vc2001 2223 8.8.8.8 mode raw
PE3(config-pseudowire)#
```

6. Attach L2-ckt under VPLS instance.

```
PE3 (config)#mpls vpls vpls2000 2000
PE3(config-vpls)#vpls-vc vc2001
PE3(config-vpls-spoke)#
```

7. Configure mac-withdrawal under VPLS instance.

```
PE3(config)#mpls vpls vpls2000 2000
PE3(config-vpls)# vpls-vc vc2001
PE3(config-vpls-spoke)# exit-spoke
PE3(config-vpls)# mac-withdrawal flush-propagate-spoke-to-mesh
PE3(config-vpls)#exit
PE3(config)#
```

Spoke with VPLS

1. Configure router LDP.

```
Spoke(config)#router ldp
Spoke(config-router)# router-id 8.8.8.8
Spoke(config-router)# transport-address ipv4 8.8.8.8
```

2. Configure targeted-peer under router LDP.

```
Spoke(config-router)# targeted-peer ipv4 5.5.5.5
```

```
Spoke(config-router-targeted-peer)# exit-targeted-peer-mode
Spoke(config-router)# targeted-peer ipv4 3.3.3.3
Spoke(config-router-targeted-peer)# exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
Spoke(config)#interface xe12
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching
```

```
Spoke(config)#interface xe25
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching
```

4. Configure VPLS instance.

```
Spoke (config)#mpls vpls vpls2000 2000
Spoke (config-vpls)#
```

5. Configure L2-ckt.

```
Spoke(config)#mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
Spoke(config-pseudowire)#!
Spoke(config-pseudowire)#mpls l2-circuit vc2001 2223 3.3.3.3 mode raw
Spoke(config-pseudowire)#
```

6. Configure Primary and Secondary spoke under VPLS instance.

```
Spoke(config)#mpls vpls vpls2000 2000
Spoke(config-vpls)#vpls-vc vc2000
Spoke(config-vpls-spoke)# secondary vc2001
Spoke(config-vpls-spoke)# exit-spoke
Spoke(config-vpls)# exit-vpls
Spoke(config)#
```

7. Configure sub-interface and attach vpls-instance to sub-interface.

```
Spoke(config)#
Spoke(config)#interface xe26.2000 switchport
Spoke(config-if)# encapsulation dot1q 2000
Spoke(config-if)# access-if-vpls
Spoke(config-acc-if-vpls)# mpls-vpls vpls2000
Spoke(config-acc-if-vpls)#
```

8. Configure mac-withdrawal under VPLS instance.

```
Spoke(config)#mpls vpls vpls2000 2000
Spoke(config-vpls)#mac-withdrawal flush-on-spoke-vc-standby-activation
Spoke(config-vpls)#commit
Spoke(config-vpls)#
```

Spoke with VPWS

1. Configure router LDP.

```
Spoke(config)#router ldp
Spoke(config-router)# router-id 8.8.8.8
Spoke(config-router)# transport-address ipv4 8.8.8.8
```

2. Configure targeted-peer under router LDP.

```
Spoke(config-router)# targeted-peer ipv4 5.5.5.5
Spoke(config-router-targeted-peer)# exit-targeted-peer-mode
Spoke(config-router)# targeted-peer ipv4 3.3.3.3
Spoke(config-router-targeted-peer)# exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
Spoke(config)#interface xe12
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching

Spoke(config)#interface xe25
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching
```

4. Configure L2-ckt.

```
Spoke(config)#mpls l2-circuit vc2222 2222 5.5.5.5 mode raw
Spoke(config-pseudowire)#!
Spoke(config-pseudowire)#mpls l2-circuit vc2223 2223 3.3.3.3 mode raw
Spoke(config-pseudowire)#
```

5. Configure sub-interface and attach VPWS-instance to sub-interface.

```
Spoke(config)#interface xe26.1222 switchport
Spoke(config-if)#encapsulation dot1q 222
Spoke(config-if)#access-if-vpws
Spoke(config-acc-if-vpls)#mpls-l2-circuit vc2222 primary
Spoke(config-acc-if-vpls)#mpls-l2-circuit vc2223 secondary
Spoke(config-acc-if-vpls)#vc-mode revertive
Spoke(config-acc-if-vpls)#mac-withdrawal flush-on-vpws-vc-switchover
```

Running Configuration

PE1

```
router ldp
router-id 2.2.2.2
targeted-peer ipv4 3.3.3.3
exit-targeted-peer-mode
targeted-peer ipv4 5.5.5.5
transport-address ipv4 2.2.2.2
!
interface xe14
enable-ldp ipv4
!
interface xe26
enable-ldp ipv4
!
mpls vpls vpls2000 2000
signaling ldp
vpls-peer 3.3.3.3
vpls-peer 5.5.5.5
exit-signaling
exit-vpls
!
interface xe16.2000 switchport
access-if-vpls
mpls-vpls vpls2000
```

PE2

```
router ldp
targeted-peer ipv4 2.2.2.2
exit-targeted-peer-mode
targeted-peer ipv4 3.3.3.3
exit-targeted-peer-mode
transport-address ipv4 5.5.5.5
!
mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
```

```

!
mpls vpls vpls2000 2000
  vpls-vc vc2000
  exit-spoke
mac-withdrawal flush-propagate-spoke-to-mesh
signaling ldp
  vpls-peer 2.2.2.2
  vpls-peer 3.3.3.3
  exit-signaling
exit-vpls

```

PE3

```

router ldp
  targeted-peer ipv4 2.2.2.2
  exit-targeted-peer-mode
  targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
transport-address ipv4 3.3.3.3
!
mpls l2-circuit vc2001 2223 8.8.8.8 mode raw
!
mpls vpls vpls2000 2000
  vpls-vc vc2001
  exit-spoke
exit-spoke
mac-withdrawal flush-propagate-spoke-to-mesh
signaling ldp
  vpls-peer 2.2.2.2
  vpls-peer 5.5.5.5
  exit-signaling
exit-vpls

```

Spoke with VPLS

```

router ldp
  router-id 8.8.8.8
  targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
  targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
transport-address ipv4 8.8.8.8
!
mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
!
mpls l2-circuit vc2001 2223 3.3.3.3 mode raw
!
mpls vpls vpls2000 2000
  vpls-vc vc2000
  secondary vc2001
  exit-spoke
mac-withdrawal flush-on-spoke-vc-standby-activation
exit-vpls
!
interface xe26.2000 switchport
access-if-vpls
mpls-vpls vpls2000

```

Spoke with VPWS

```

router ldp

```



```

router-id 8.8.8.8
targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
transport-address ipv4 8.8.8.8
!
mpls l2-circuit vc2222 2222 5.5.5.5 mode raw
!
mpls l2-circuit vc2223 2223 3.3.3.3 mode raw
!
interface xe26.1222 switchport
encapsulation dot1q 222
access-if-vpws
  mpls-l2-circuit vc2222 primary
  mpls-l2-circuit vc2223 secondary
  vc-mode revertive
  mac-withdrawal flush-on-vpws-vc-switchover
!

```

Validation

PE2#show mpls vpls vpls2000

```

Virtual Private LAN Service Instance: vpls2000, ID: 2000
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
MAC Withdrawal:
  Propagated to the mesh peers from the hub

Configured interfaces:
  None

Mesh Peers:
  2.2.2.2 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h13m)
  3.3.3.3 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h22m)
Spoke Peers:
  vc2000 (Up) (UpTime 01:31:27)

```

PE3#show mpls vpls vpls2000

```

Virtual Private LAN Service Instance: vpls2000, ID: 2000
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
MAC Withdrawal:
  Propagated to the mesh peers from the hub

Configured interfaces:
  None

```

```

Mesh Peers:
  2.2.2.2 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h16m)
  5.5.5.5 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d12h22m)
Spoke Peers:
  vc2001 (Dn) (Reason: VC on standby)

```

For VPLS: Spoke#show mpls l2-circuit detail

```

Spoke#show mpls l2-circuit detail
Virtual Private LAN Service Instance: vpls2000, ID: 2000
SIG-Protocol: N/A
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 1500
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:
Sent on switchover from secondary spoke to primary spoke
Configured interfaces:
Interface: xe26.2000
Status: Up
Subinterface Match Criteria(s) :
dot1q 2000
Spoke Peers:
vc2000 (Up) (UpTime 01:31:33)
Secondary: vc2001 (Dn) (Reason: VC on standby)

```

For VPWS: Spoke#show mpls l2-circuit detail

```

spoke#show mpls l2-circuit detail
MPLS Layer-2 Virtual Circuit: vc2222, id: 3028 PW-INDEX: 1 service-tpid: dot1.
q
Endpoint: 5.5.5.5
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
In-label: 24960, Out-label: 24962
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe26.1222
  Subinterface Match Criteria(s) :
    dot1q 222

Virtual Circuit Type: Ethernet VLAN
Virtual Circuit is configured as Primary
Virtual Circuit is configured as Revertive
Local PW Status :
  Forwarding
Remote PW Status :
  Forwarding
Local VCCV Capability:
CC-Types:  Type 2(in use)
CV-Types: LSP ping(in use)

Virtual Circuit is active (State: Installed)
  CreateTime: 00:29:32, UpTime: 00:04:22, FlapCount: 1, LastFlapReason: AC Inte
rface down

MPLS Layer-2 Virtual Circuit: vc2223, id: 3032 PW-INDEX: 2 service-tpid: dot1.
q
Endpoint: 3.3.3.3
Control Word: 0
Flow Label Status: Disabled, Direction: None, Static: No
In-label: 24961, Out-label: 24962
MPLS Layer-2 Virtual Circuit Group: none
Bound to interface: xe26.1222
  Subinterface Match Criteria(s) :

```

```
dot1q 222
```

```
Virtual Circuit Type: Ethernet VLAN  
Virtual Circuit is configured as Secondary  
Virtual Circuit is configured as Revertive  
Local PW Status :
```

```
  PW Standby
```

```
Remote PW Status :  
  Forwarding
```

```
Local VCCV Capability:
```

```
CC-Types:  Type 2(in use)
```

```
CV-Types:  LSP ping(in use)
```

```
Virtual Circuit is inactive (State: Signaled)
```

```
  CreateTime: 00:29:32, DownReason: Primary entity UP
```

CLI Commands

The MAC Withdrawal for H-VPLS with VPWS introduces the following configuration commands.

mac-withdrawal flush-on-spoke-vc-standby-activation

Use this command to send MAC withdraw when the secondary spoke becomes the primary.

Use `no` parameter of this command to remove MAC withdraw when the secondary spoke becomes the primary.

Command Syntax

```
mac-withdrawal flush-on-spoke-vc-standby-activation
```

```
no mac-withdrawal flush-on-spoke-vc-standby-activation
```

Parameters

None

Default

None

Command Mode

VPLS Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

The following example is for enabling MAC Withdrawal on Spoke VC Activation:

```
#configure terminal
(config)#mpls vpls vpls12 12
(config-vpls)#mac-withdrawal flush-on-spoke-vc-standby-activation
```

mac-withdrawal flush-propagate-spoke-to-mesh

Use this command to allow propagation of MAC withdraw from spoke to mesh.

Use `no` parameter of this command to remove propagation of MAC withdraw from spoke to mesh.

Command Syntax

```
mac-withdrawal flush-propagate-spoke-to-mesh
```

```
no mac-withdrawal flush-propagate-spoke-to-mesh
```

Parameters

None

Default

None

Command Mode

VPLS Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Explain or describe the example.

```
#configure terminal
(config)#mpls vpls vpls12 12
(config-vpls)#mac-withdrawal flush-propagate-spoke-to-mesh
```

mac-withdrawal flush-on-vpws-vc-switchover

Use this command to send MAC withdraw when the secondary spoke becomes the primary and primary becomes active due to vc mode revertive.

Use `no` parameter of this command to remove MAC withdraw when the secondary spoke becomes the primary.

Command Syntax

```
mac-withdrawal flush-on-vpws-vc-switchover
```

```
no mac-withdrawal flush-on-vpws-vc-switchover
```

Parameters

None

Default

None

Command Mode

MPLS Mode

Applicability

Introduced in OcNOS version 6.6.1.

Example

The following example is for

```
(config)#interface xe26.1222 switchport
(config-if)#encapsulation dot1q 222
(config-if)#access-if-vpws
(config-acc-if-vpws)#mpls-l2-circuit vc-2222 primary
(config-acc-if-vpws)#mpls-l2-circuit vc-2223 secondary
(config-acc-if-vpws)#vc-mode revertive
(config-acc-if-vpws)#mac-withdrawal flush-on-vpws-vc-switchover
(config-acc-if-vpws)#
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms/Acronym | Description |
|-------------------|--|
| H-VPLS | Hierarchical Virtual Private LAN Service |
| VPLS | Virtual Private LAN Service |

MAC Move Protection - VPLS/H-VPLS

Overview

MAC Move Protection is a Layer 2 feature for detecting and managing the movement of MAC addresses across various interfaces in Virtual Private LAN Service (VPLS) or Hierarchical VPLS (H-VPLS) networks.

In VPLS environments, MAC address moves can occur across Attachment Circuits (AC), Spoke-PWs, and Mesh-PWs. MAC Move Protection is particularly useful in detecting and responding to these movements within these different components.

Characteristics of MAC Move Protection - VPLS/H-VPLS

- Monitors MAC address movements across Attachment Circuits (AC), Spoke-PWs, and Mesh-PWs, detecting any moves between these components in a VPLS/H-VPLS topology.
- Enables detection settings across multiple VPLS instances, ensuring uniformity and reducing redundant configurations.
- Allows more granular control, enabling overrides for detection timers and error-disable actions on specific instances.
- Administrators can configure detection interval and move count threshold.
- When a MAC move is detected, this feature:
 - Applies error-disable actions to ACs to prevent disruption.
 - Brings down Spoke-PWs or Mesh-PWs, reducing impact on the network.
 - Automatically restores components after the error-disable or operational down actions are triggered.

Benefits

- By detecting and managing unexpected MAC address moves, the feature helps prevent network loops, service disruptions, and performance degradation, ensuring stable VPLS connectivity.
- The action mechanism minimizes disruptions by intelligently deciding which components to block (AC, Spoke-PW, or Mesh-PW) based on priority, reducing the impact of MAC move events on the overall network.
- With syslog reporting and detailed CLI commands, network administrators can quickly identify and address MAC move issues.
- The ability to configure detection settings both globally and at the instance level provides flexibility in managing large-scale VPLS networks.

Prerequisites

- Define Interfaces and Loopback Addresses:

Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

!


```

interface lo
  ip address 127.0.0.1/8
  ip address 2.2.2.2/32 secondary
  ipv6 address ::1/128

interface xe14
  ip address 30.1.1.2/24

```

- **Configure IGP for Dynamic Routing:**

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

- **ISIS Configuration:**

```

router isis 1
  is-type level-2-only
  metric-style wide
  microloop-avoidance level-2
  mpls traffic-eng router-id 2.2.2.2
  mpls traffic-eng level-2
  capability cspf
  dynamic-hostname
    bfd all-interfaces
  net 49.0000.0000.0002.00
  passive-interface lo
  !
interface xe14
  mpls ldp-igp sync isis level-2
  isis network point-to-point
  ip router isis 1

```

- **OSPF Configuration:**

```

router ospf 1
  ospf router-id 2.2.2.2
  network 2.2.2.2/32 area 0.0.0.0
  network 30.1.1.0/24 area 0.0.0.0!
  !
interface xe14
  ip ospf network point-to-point


```

Configuration

Topology

The sample topology for MAC Move Protection with CE-PE-Hub-Spoke architecture.

Figure 83. MAC Move Protection Topology Diagram



PE1

1. Configure router LDP.

```
PE1(config)#router ldp
PE1(config-router)# router-id 2.2.2.2
PE1(config-router)# transport-address ipv4 2.2.2.2
```

2. Configure targeted-peer under router LDP.

```
PE1(config-router)# targeted-peer ipv4 5.5.5.5
PE1(config-router-targeted-peer)# exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
PE1(config)#interface xe14
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching
```

4. Configure VPLS instance.

```
PE1(config)#mpls vpls vpls2000 2000
PE1(config-vpls)# signaling ldp
PE1(config-vpls-sig)# vpls-peer 5.5.5.5
PE1(config-vpls-sig)# exit-signaling
PE1(config-vpls)# exit-vpls
PE1(config)#
```

5. Configure sub-interface and attach vpls-instance to sub-interface.

```
PE1(config)#
PE1(config)#interface xe16.2000 switchport
PE1(config-if)# encapsulation dot1q 2000
PE1(config-if)# access-if-vpls
PE1(config-acc-if-vpls)# mpls-vpls vpls2000
PE1(config-acc-if-vpls)#
```

6. MAC-MOVE global configuration.

```
PE1(config)#vpls mac-move enable detect 10 10
PE1(config)#commit
PE1(config)#
```

7. Configure MAC-MOVE under VPLS instance.

```
PE1(config)#mpls vpls vpls2000 2000
PE1(config-vpls)# mac-move
PE1(config-vpls-mac-move)# detect 10 10
PE1(config-vpls-mac-move)# errdisable timeout-interval 120
PE1(config-vpls-mac-move)#commit
PE1(config-vpls-mac-move)#
PE1(config-vpls-mac-move)#exit
PE1(config)#
```

Hub

1. Configure router LDP.

```
Hub(config)#router ldp
Hub(config-router)# router-id 5.5.5.5
Hub(config-router)# transport-address ipv4 5.5.5.5
```

2. Configure targeted-peer under router LDP.

```
Hub(config-router)# targeted-peer ipv4 2.2.2.2
Hub(config-router-targeted-peer)# exit-targeted-peer-mode
R5-P5(config-router)# targeted-peer ipv4 8.8.8.8
R5-P5(config-router-targeted-peer)#
```

3. Enable LDP and label-switching for core interface.

```
Hub(config)#interface xe1
Hub(config-if)# enable-ldp ipv4
Hub(config-if)#label-switching
Hub(config)#interface xe12
Hub(config-if)# enable-ldp ipv4
Hub(config-if)#label-switching
```

4. Configure VPLS instance.

```
Hub(config)#mpls vpls vpls2000 2000
Hub(config-vpls)# signaling ldp
Hub(config-vpls-sig)# vpls-peer 2.2.2.2
Hub(config-vpls-sig)# exit-signaling
Hub(config-vpls)# exit-vpls
Hub(config)#
```

5. Configure L2-ckt.

```
Hub (config)#mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
Hub (config-pseudowire)#
```

6. Attach L2-ckt under VPLS instance.

```
Hub (config)#mpls vpls vpls2000 2000
Hub (config-vpls)#vpls-vc vc2000
Hub (config-vpls-spoke)#
```

7. MAC Move Protection global configuration.

```
Hub(config)#vpls mac-move enable detect 10 10
Hub(config)#commit
Hub(config)#
```

8. Configure MAC-MOVE under VPLS instance.

```
Hub(config)#mpls vpls vpls2000 2000
Hub(config-vpls)# mac-move
Hub(config-vpls-mac-move)# detect 10 10
Hub(config-vpls-mac-move)# errdisable timeout-interval 120
Hub(config-vpls-mac-move)# errdisable allow-mesh-pw-blocking
Hub(config-vpls-mac-move)#exit
Hub(config)#PE1(config)#
```

Spoke

1. Configure router LDP.

```
Spoke(config)#router ldp
Spoke(config-router)# router-id 8.8.8.8
Spoke(config-router)# transport-address ipv4 8.8.8.8
```

2. Configure targeted-peer under router LDP.

```
Spoke(config-router)# targeted-peer ipv4 5.5.5.5
Spoke(config-router-targeted-peer)# exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
Spoke(config)#interface xe12
Spoke(config-if)# enable-ldp ipv4
Spoke(config-if)#label-switching
```

4. Configure VPLS instance.

```
Spoke(config)#mpls vpls vpls2000 2000
Spoke(config-vpls)#
```

5. Configure L2-ckt.

```
Spoke(config)#mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
Spoke(config-pseudowire)#
```

6. Attach L2-ckt under VPLS instance.

```
Spoke (config)#mpls vpls vpls2000 2000
Spoke(config-vpls)#vpls-vc vc2000
Spoke(config-vpls-spoke)#
```

7. MAC Move Protection global configuration

```
Spoke(config)#vpls mac-move enable detect 10 10
Spoke(config)#commit
Spoke(config)#
```

8. Configure MAC-MOVE under VPLS instance.

```
Spoke(config)#mpls vpls vpls2000 2000
Spoke(config-vpls)# mac-move
Spoke(config-vpls-mac-move)# detect 10 10
Spoke(config-vpls-mac-move)# errdisable timeout-interval 120
Spoke(config-vpls-mac-move)# errdisable allow-mesh-pw-blocking
Spoke(config-vpls-mac-move)#exit
Spoke(config)#PE1(config)#
```

Running Configuration

PE1

```
router ldp
router-id 2.2.2.2
targeted-peer ipv4 5.5.5.5
exit-targeted-peer-mode
transport-address ipv4 2.2.2.2
!
interface xe14
enable-ldp ipv4
!
mpls vpls vpls2000 2000
signaling ldp
vpls-peer 5.5.5.5
exit-signaling
mac-move
detect 10 10
errdisable timeout-interval 120
exit-mac-move
exit-vpls
!
interface xe16.2000 switchport
access-if-vpls
mpls-vpls vpls2000
learning limit prof1
```

Hub

```
router ldp
targeted-peer ipv4 2.2.2.2
exit-targeted-peer-mode
targeted-peer ipv4 8.8.8.8
exit-targeted-peer-mode
!
!
```

```

mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
!
mpls vpls vpls2000 2000
  vpls-vc vc2000
  learning limit prof1
  exit-spoke
  signaling ldp
  vpls-peer 2.2.2.2
  exit-signaling
mac-move
  detect 10 10
  errdisable timeout-interval 120
  errdisable allow-mesh-pw-blocking
  exit-mac-move
exit-vpls
exit-vpls

```

Spoke

```

router ldp
  router-id 8.8.8.8
  targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
  transport-address ipv4 8.8.8.8
!
mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
!
mpls vpls vpls2000 2000
  vpls-vc vc2000
  exit-spoke
  exit-signaling
mac-move
  detect 10 10
  errdisable timeout-interval 120
  errdisable allow-mesh-pw-blocking
  exit-mac-move
exit-vpls!
interface xe26.2000 switchport
  access-if-vpls
  mpls-vpls vpls2000

```

Validation

When mac move is seen on Hub:

```

-----
HUB#2025 Jan 22 11:12:34.684 : HUB : NSM : NOTIF : [IFMGR_ERR_DISABLE_UP_4]: Mesh with Peer 2.2.2.2
on VPLS instance vpls2000 recovered from operational shutdown
2025 Jan 22 11:12:34.687 : HUB : NSM : NOTIF : [NSM_MPLS_VPLS_PEER_STATE_CHANGE_4]: VPLS vpls2000 ID
2000 peer 2.2.2.2 changed state to up
2025 Jan 22 11:12:34.695 : HUB : NSM : CRITI : [IFMGR_ERR_DISABLE_DOWN_2]: Mesh with peer 3.3.3.3 on
VPLS instance vpls2000 shutdown successfully
2025 Jan 22 11:12:34.697 : HUB : NSM : CRITI : [NSM_MPLS_VPLS_PEER_STATE_CHANGE_2]: VPLS vpls2000 ID
2000 peer 3.3.3.3 changed state to down (Reason: VPLS peer errdisable)
2025 Jan 22 11:12:37.196 : HUB : HSL : CRITI : L2 movement detected 221 times : sample MAC :
0000:0000:0009 from PEER : 2.2.2.2

```

Hub#show mpls vpls vpls2000

```

Virtual Private LAN Service Instance: vpls2000, ID: 2000
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Enabled

```

```

Flow Label Status: Enabled, Direction: Both, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 5000
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe2.2000
Status: Up
Subinterface Match Criteria(s) :
dot1q 2000

Mesh Peers:
2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: Yes, FAT: No) (Up) (UpTime: 2d00h01m)
FEC signaling element: FEC128
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: Yes, FAT: No) (Up) (UpTime: 01:44:45)
FEC signaling element: FEC128
Spoke Peers:
vc2000 (Dn) (Reason: VPLS peer errdisable)

```

When mac move is cleared on Hub:

```

-----
HUB#2025 Jan 22 11:17:34.697 : HUB : NSM : NOTIF : [IFMGR_ERR_DISABLE_UP_4]: Mesh with Peer 3.3.3.3
on VPLS instance vpls2000 recovered from operational shutdown
2025 Jan 22 11:17:34.700 : HUB : NSM : NOTIF : [NSM_MPLS_VPLS_PEER_STATE_CHANGE_4]: VPLS vpls2000 ID
2000 peer 3.3.3.3 changed state to up

```

Hub#show mpls vpls vpls2000

```

Virtual Private LAN Service Instance: vpls2000, ID: 2000
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Enabled
Flow Label Status: Enabled, Direction: Both, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 5000
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe2.2000
Status: Up
Subinterface Match Criteria(s) :
dot1q 2000

Mesh Peers:
2.2.2.2 (Type: Ethernet VLAN) (Negotiated - CW: Yes, FAT: No) (Up) (UpTime: 2d00h01m)
FEC signaling element: FEC128
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: Yes, FAT: No) (Up) (UpTime: 01:44:45)
FEC signaling element: FEC128

Spoke Peers:
vc2000 (Dn) (Reason: VC on standby)

```

When mac move is seen on PE1:

```

-----

```

PE1#show mpls vpls vpls2001

```
Virtual Private LAN Service Instance: vpls2001, ID: 2001
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Enabled
Flow Label Status: Enabled, Direction: Both, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 5000
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe16.2001
Status: Down
Subinterface Match Criteria(s) :
dot1q 2001

Mesh Peers:
3.3.3.3 (Type: Ethernet VLAN) (Negotiated - CW: Yes, FAT: Yes) (Up) (UpTime: 01:53:26)
FEC signaling element: FEC128
5.5.5.5 (Type: Ethernet VLAN) (Negotiated - CW: Yes, FAT: Yes) (Up) (UpTime: 2d00h09m)
FEC signaling element: FEC128
```

PE1#show interface brief | grep xe16.2001

| | | | | | | | | | |
|-----------|--------------|----|----|------|----|-----|----|----|----|
| xe16.2001 | SUBINTERFACE | -- | -- | down | ED | 10g | -- | No | No |
| PE1# | | | | | | | | | |

CLI Commands

The MAC Move Protection introduces the following configuration commands.

vpls mac-move enable detect

Use this command to enable MAC address move detection within a VPLS environment with global configuration.

Use `no` parameter of this command to disable MAC address move detection

Command Syntax

```
vpls mac-move enable detect <1-1000> <5-300>
```

```
no vpls mac-move enable detect
```

Parameters

<1-1000>

Specifies the number of detected MAC address moves required to trigger an action. The default value is 5.

<5-300>

Specifies the time period (in seconds) within which the specified number of MAC address moves must occur for the move to be considered valid. The default value is 15 seconds.

Default

Disabled

Command Mode

CONFIG mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

The following example is for configuration of MAC move protection using global configuration:

```
#configure terminal
(config)#vpls mac-move enable detect 10 40
(config)#commit
```

mac-move

Use this command to enable MAC address move detection within a VPLS environment with VPLS MAC MOVE mode.

Use `no` parameter of this command to disable MAC address move detection

Command Syntax

```
mac-move detect (<1-1000> | <5-300>) | errdisable (allow-mesh-pw-blocking | timeout-interval <0-86400>)
```

```
no mac-move
```

Parameters

<1-1000>

Specifies the number of detected MAC address moves required to trigger an action. The default value is 5.

<5-300>

Specifies the time period (in seconds) within which the specified number of MAC address moves must occur for the move to be considered valid. The default value is 15 seconds.

allow-mesh-pw-blocking

Allows blocking the Mesh Pseudowire (PW) instead of only disabling the MAC in case of an error.

<0-86400>

(Optional) Specifies the MAC move errdisable timeout interval, determining how long the affected MAC remains disabled before being re-enabled. The default value is 0 second.

Default

Disabled

Command Mode

VPLS MAC MOVE mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

The following example is for configuration of MAC move protection for VPLS instance:

```
#configure terminal
(config)# mpls vpls vpls_test1 100
(config-vpls)#mac-move
(config-vpls-mac-move)# detect 10 60
(config-vpls-mac-move)# errdisable timeout interval 120
(config-vpls-mac-move)# errdisable allow-mesh-pw-blocking
(config-vpls-mac-move)# exit-mac-move
(config-vpls)# exit-vpls
```

show mpls vpls mac-move name

Use this command to display the MAC address move configuration and status for the VPLS instance.

Command Syntax

```
show mpls vpls mac-move name
```

Parameters

None

Applicability

Introduced in OcNOS version 6.6.0.

Example

The following example is for configuration of MAC move protection:

```
#show mpls vpls mac-move name vpls26
Virtual Private LAN Service Instance: vpls26, ID:26
Mac Address      Move Count    Elapsed time
90:67:17:e2:46:74  29             00:17:35
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms/Acronym | Description |
|-------------------|--|
| CLI | Command Line Interface |
| H-VPLS | Hierarchical Virtual Private LAN Service |
| IGP | Interior Gateway Protocol |
| ISIS | Intermediate System to Intermediate System |
| OSPF | Open Shortest Path First |
| BFD | Bidirectional Forwarding Detection |
| VPLS | Virtual Private LAN Service |

MAC Limit for VPLS and H-VPLS

Overview

The MAC limit controls how many MAC addresses a system can learn, which is especially beneficial in Virtual Private LAN Service (VPLS) deployments. This control allows you to limit MAC addresses at more granular levels, such as the Access Circuit (AC) or Spoke-PW level, while maintaining the current VPLS instance-level limits.

Characteristics of MAC Move Limit - VPLS/H-VPLS

- Threshold-based control for the number of MAC addresses.
- Granular configuration options, including interface/subinterface/AC/Spoke-PW levels.
- Monitoring and enforcement with actions like logging or error-disable.
- Security benefits by preventing MAC flooding and limiting device access.
- Non-disruptive operation with logging, and optional error-disable with recovery options.

Benefits

- Prevents MAC flooding attacks, controls access to network segments.
- Improves network efficiency by managing memory and CPU usage.
- Granular configuration at interface, subinterface, AC, and Spoke-PW levels.
- Prevents MAC table overflows, ensuring stable traffic forwarding.
- Syslog alerts and watermark thresholds for proactive management.
- Logging doesn't affect traffic, and error-disable includes recovery options.
- Helps networks grow efficiently without overloading devices.

Prerequisites

- Define Interfaces and Loopback Addresses:

Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 2.2.2.2/32 secondary  
  ipv6 address ::1/128  
  
interface xe14  
  ip address 30.1.1.2/24
```

- Configure IGP for Dynamic Routing:

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal

network performance.

- **ISIS Configuration:**

```
router isis 1
 is-type level-2-only
 metric-style wide
 microloop-avoidance level-2
 mpls traffic-eng router-id 2.2.2.2
 mpls traffic-eng level-2
 capability cspf
 dynamic-hostname
   bfd all-interfaces
 net 49.0000.0000.0002.00
 passive-interface lo
 !
 interface xe14
 mpls ldp-igp sync isis level-2
 isis network point-to-point
 ip router isis 1
```

- **OSPF Configuration:**


```
router ospf 1
 ospf router-id 2.2.2.2
 network 2.2.2.2/32 area 0.0.0.0
 network 30.1.1.0/24 area 0.0.0.0!
 !
 interface xe14
 ip ospf network point-to-point
```

Configuration

Topology

The sample topology for MAC Limit with CE-PE-Hub-Spoke architecture.

Figure 84. MAC Limit for H-VPLS Topology Diagram



PE1

1. Configure router LDP.

```
PE1(config)#router ldp
PE1(config-router)# router-id 2.2.2.2
PE1(config-router)# transport-address ipv4 2.2.2.2
```

2. Configure targeted-peer under router LDP.

```
PE1(config-router)# targeted-peer ipv4 5.5.5.5
PE1(config-router-targeted-peer)# exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
PE1(config)#interface xe14
```

```
PE1(config-if)# enable-ldp ipv4
PE1(config-if)# label-switching
```

4. Configure VPLS instance.

```
PE1(config)# mpls vpls vpls2000 2000
PE1(config-vpls)# signaling ldp
PE1(config-vpls-sig)# vpls-peer 5.5.5.5
PE1(config-vpls-sig)# exit-signaling
PE1(config-vpls)# exit-vpls
PE1(config)#
```

5. Configure sub-interface and attach vpls-instance to sub-interface.

```
PE1(config)#
PE1(config)# interface xe16.2000 switchport
PE1(config-if)# encapsulation dot1q 2000
PE1(config-if)# access-if-vpls
PE1(config-acc-if-vpls)# mpls-vpls vpls2000
PE1(config-acc-if-vpls)#
```

6. Configure MAC Limit profile configuration.

```
PE1(config)#
PE1(config)# vpls mac-limit-profile prof1
PE1(config-vpls-mac-lim-profile)# learning-limit 5
PE1(config-vpls-mac-lim-profile)# action log-errdisable
PE1(config-vpls-mac-lim-profile)# errdisable-timeout 120
```

7. Configure MAC Limit profile under AC sub-interface.

```
PE1(config)#
PE1(config)# interface xe16.2001 switchport
PE1(config-if)# access-if-vpls
PE1(config-acc-if-vpls)# learning limit prof1
PE1(config-acc-if-vpls)# exit
```

Hub

1. Configure router LDP.

```
Hub(config)# router ldp
Hub(config-router)# router-id 5.5.5.5
Hub(config-router)# transport-address ipv4 5.5.5.5
```

2. Configure targeted-peer under router LDP.

```
Hub(config-router)# targeted-peer ipv4 2.2.2.2
Hub(config-router-targeted-peer)# exit-targeted-peer-mode
Hub(config-router)# targeted-peer ipv4 8.8.8.8
Hub(config-router-targeted-peer)#
```

3. Enable LDP and label-switching for core interface.

```
Hub(config)# interface xe1
Hub(config-if)# enable-ldp ipv4
Hub(config-if)# label-switching
Hub(config)# interface xe12
Hub(config-if)# enable-ldp ipv4
Hub(config-if)# label-switching
```

4. Configure VPLS instance.

```
Hub(config)# mpls vpls vpls2000 2000
Hub(config-vpls)# signaling ldp
Hub(config-vpls-sig)# vpls-peer 2.2.2.2
Hub(config-vpls-sig)# exit-signaling
Hub(config-vpls)# exit-vpls
Hub(config)#
```

5. Configure L2-ckt.

```
Hub (config)#mpls l2-circuit vc2000 2222 8.8.8.8 mode raw
Hub (config-pseudowire)#
```

6. Attach L2-ckt under vpls instance.

```
Hub (config)#mpls vpls vpls2000 2000
Hub (config-vpls)#vpls-vc vc2000
Hub (config-vpls-spoke)#
```

7. Configure MAC Limit profile configuration.

```
HUB(config)#vpls mac-limit-profile prof1
HUB(config-vpls-mac-lim-profile)# learning-limit 5
HUB(config-vpls-mac-lim-profile)# action log-errdisable
HUB(config-vpls-mac-lim-profile)# errdisable-timeout 120
HUB(config-vpls-mac-lim-profile)#
```

8. Configure MAC Limit profile under VPLS instance.

```
HUB(config)#mpls vpls vpls2001 2001
HUB(config-vpls)#vpls-vc vc2000
HUB(config-vpls-spoke)# learning limit prof1
HUB(config-vpls-spoke)#
```

Spoke

1. Configure router LDP.

```
Spoke (config)#router ldp
Spoke (config-router)# router-id 8.8.8.8
Spoke (config-router)# transport-address ipv4 8.8.8.8
```

2. Configure targeted-peer under router LDP.

```
Spoke (config-router)# targeted-peer ipv4 5.5.5.5
Spoke (config-router-targeted-peer)# exit-targeted-peer-mode
```

3. Enable LDP and label-switching for core interface.

```
Spoke (config)#interface xe12
Spoke (config-if)# enable-ldp ipv4
Spoke (config-if)#label-switching
```

4. Configure VPLS instance.

```
Spoke (config)#mpls vpls vpls2000 2000
Spoke (config-vpls)#
```

5. Configure L2-ckt.

```
Spoke (config)#mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
Spoke (config-pseudowire)#
```

6. Attach L2-ckt under VPLS instance.

```
Spoke (config)#mpls vpls vpls2000 2000
Spoke (config-vpls)#vpls-vc vc2000
Spoke (config-vpls-spoke)#
```

7. Configure sub-interface and attach vpls-instance to sub-interface.

```
Spoke (config)#
Spoke (config)#interface xe26.2000 switchport
Spoke (config-if)# encapsulation dot1q 2000
Spoke (config-if)# access-if-vpls
Spoke (config-acc-if-vpls)# mpls-vpls vpls2000
Spoke (config-acc-if-vpls)#
```

8. Configure MAC Limit profile configuration.

```
Spoke(config)#vpls mac-limit-profile R8
Spoke(config-vpls-mac-lim-profile)# learning-limit 10
Spoke(config-vpls-mac-lim-profile)# action log-errdisable
Spoke(config-vpls-mac-lim-profile)# errdisable-timeout 60
Spoke(config-vpls-mac-lim-profile)#
Spoke(config-vpls-mac-lim-profile)#
```

9. Configure MAC Limit profile under vpls instance.

```
Spoke(config)#mpls vpls vpls2000 2000
Spoke(config-vpls)#vpls-vc vc2000
Spoke(config-vpls-spoke)#learning limit R8
Spoke(config-vpls-spoke)#
```

Running Configuration**PE1**

```
vpls mac-limit-profile prof1
  learning-limit 5
  action log-errdisable
  errdisable-timeout 120
!

router ldp
  router-id 2.2.2.2
  targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
  transport-address ipv4 2.2.2.2
!
interface xe14
  enable-ldp ipv4
!
mpls vpls vpls2000 2000
  signaling ldp
  vpls-peer 5.5.5.5
  exit-signaling
  exit-vpls
!
interface xe16.2000 switchport
  access-if-vpls
  mpls-vpls vpls2000
  learning limit prof1
```

Hub

```
vpls mac-limit-profile prof1
  learning-limit 5
  action log-errdisable
  errdisable-timeout 120
!

router ldp
  targeted-peer ipv4 2.2.2.2
  exit-targeted-peer-mode
  targeted-peer ipv4 8.8.8.8
  exit-targeted-peer-mode
!
!
mpls 12-circuit vc2000 2222 8.8.8.8 mode raw
!
```



```

mpls vpls vpls2000 2000
  vpls-vc vc2000
  learning limit prof1
  exit-spoke
  signaling ldp
  vpls-peer 2.2.2.2
  exit-signaling
  exit-vpls

```

Spoke

```

vpls mac-limit-profile R8
  learning-limit 10
  action log-errdisable
  errdisable-timeout 60
!

router ldp
  router-id 8.8.8.8
  targeted-peer ipv4 5.5.5.5
  exit-targeted-peer-mode
  transport-address ipv4 8.8.8.8
!
mpls l2-circuit vc2000 2222 5.5.5.5 mode raw
!
mpls vpls vpls2000 2000
  vpls-vc vc2000
  learning limit R8
  exit-spoke
  exit-vpls
!
interface xe26.2000 switchport
  access-if-vpls
  mpls-vpls vpls2000

```

Validation

Verify VPLS mesh are up between PE and Hub

```

PE1#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | | | | | |
| 2000 | 5.5.5.5 | 31364 | 28162 | xe14 | 26883 | 2/Up | 4 |
| LDP | Active | 2d10h36m | | | | | |

```

Hub#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | | | | | |
| 2000 | 2.2.2.2 | 29446 | 26883 | xe1 | 28162 | 2/Up | 3 |
| LDP | Active | 2d10h39m | | | | | |

Verify VPLS spoke are up between Hub and Spoke

```

Hub#show ldp mpls-l2-circuit
Transport Client VC Local Remote Destination Lo-
cal
VC ID Binding State Type VC Label VC Label Address PW
Status PW Status

```

```

2222          VPLS:2000  UP          Ethernet      26882      26886      8.8.8.8      Forwarding
          Forwarding

Hub#sho mpls vpls spoke
VPLS-ID      Virtual Circuit  Tunnel-Label In-Label    Network-Intf Out-Label    Lkps/St      Secondary
2000         vc2000          29443        26882       ce4          26886        2/Up
---

Spoke#show ldp mpls-l2-circuit
Transport    Client      VC          VC          Local      Remote      Destination    Lo-
cal
VC ID        Binding    State      Type          VC Label    VC Label    Address          PW
Status
2222         VPLS:2000  UP          Ethernet      26886       26882       5.5.5.5          Forwarding
          Forwarding

Spoke#show mpls vpls spoke
VPLS-ID      Virtual Circuit  Tunnel-Label In-Label    Network-Intf Out-Label    Lkps/St      Secondary
2000         vc2000          29440        26886       ce4          26882        2/Up          ---

```

Verify MAC-LIMIT session on Hub and spoke:

```

Hub#show mpls vpls vpls2000
Virtual Private LAN Service Instance: vpls2000, ID: 2000
SIG-Protocol: LDP
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet, Configured MTU: 1500
Description: none
service-tpid: dot1.q
Operating mode: Raw
Ignoring AC interface and spoke-VC state

Configured interfaces:
  None

Mesh Peers:
  2.2.2.2 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d10h47m)
  3.3.3.3 (Peer VPLS Type: Ethernet) (Up) (UpTime: 2d10h56m)
Spoke Peers:
  vc2000 (Up) (UpTime 00:05:48)

```

CLI Commands

The MAC Limit introduces the following configuration commands.

vpls mac-limit-profile

Use this command to set the MAC address learning limits which will be used to associate the AC or Spoke PW for a specific VPLS MAC limit profile.

Use `no` parameter of this command to delete the VPLS MAC limit profile.

Command Syntax

```
vpls mac-limit-profile <PROFILE_NAME> learning-limit <1-32767> | high-watermark <1-100> | low-watermark <1-100> | action (log-errdisable <0-86400> | log-only) | errdisable-timeout <0-86400>
```

```
no vpls mac-limit-profile
```

Parameters

<PROFILE_NAME>

Specifies the name of the MAC limit profile.

learning-limit <1-32767>

Specifies the maximum number of MAC addresses allowed to be learned on the interface. The default value is 32767.

high-watermark <1-100>

Specifies the high watermark (maximum number of MAC addresses) for logging purposes. The threshold is a numeric value and a percentage of the learning limit. The default value is 90%.

low-watermark <1-100>

Specifies the low watermark (minimum number of MAC addresses) for logging purposes. The threshold is a numeric value and a percentage of the learning limit. The default value is 70%.

action log-errdisable <0-86400>

Logs an event when the MAC limit is exceeded and disables MAC learning for the timeout period. The default value is 0.

action log-only

Logs when the MAC limit is exceeded without disabling MAC learning.

errdisable-timeout <0-86400>

Specifies the duration (in seconds) before MAC learning is re-enabled after being errdisabled. The default value is 0, meaning no automatic recovery.

Default

None

Command Mode

VPLS MAC Limit Profile Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

The following example is for creating a VPLS MAC limit profile and configuring with specific parameters to manage MAC address learning limits:

```
#configure terminal
(config)#vpls mac-limit-profile prof1
(config-vpls-mac-lim-profile)#learning-limit 50
(config-vpls-mac-lim-profile)#action log-errdisable
(config-vpls-mac-lim-profile)#high-watermark 60
(config-vpls-mac-lim-profile)#low-watermark 30
(config-vpls-mac-lim-profile)#errdisable-timeout 30
(config-vpls-mac-lim-profile)#commit
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:


| Key Terms/Acronym | Description |
|-------------------|--|
| CLI | Command Line Interface |
| H-VPLS | Hierarchical Virtual Private LAN Service |
| IGP | Interior Gateway Protocol |
| ISIS | Intermediate System to Intermediate System |
| OSPF | Open Shortest Path First |
| BFD | Bidirectional Forwarding Detection |
| VPLS | Virtual Private LAN Service |

MLAG Active-Standby for VPLS

Overview

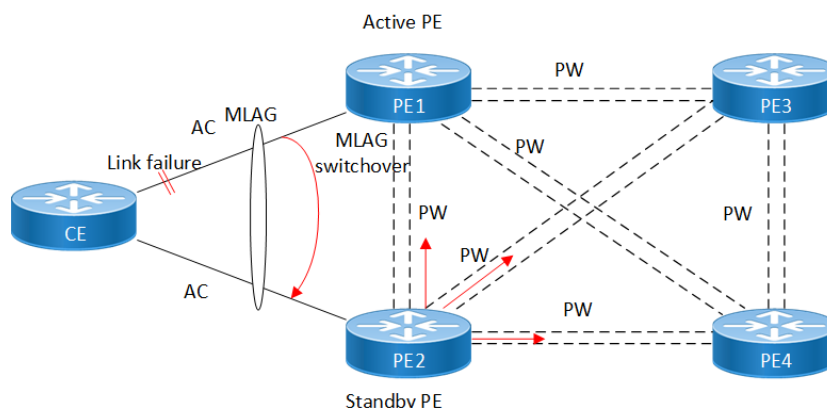
The Virtual Private LAN Service (VPLS) architecture provides a multipoint ethernet connection using the Multi-Protocol Label Switching (MPLS) transport. This helps connect multiple devices from different geographical locations to the same virtual network. The Multi-Chassis Link Aggregation (MLAG) provides the redundancy technique in the VPLS architecture, making it fail-safe multipoint VPLS connectivity. The redundancy in the network traffic is achieved by dual-homing a Customer Edge (CE) device to two Provider Edge (PE) devices.

A single CE is dual-homed to two PEs for redundancy is shown in the following figure:



Feature Characteristics

The Multi-Chassis Link Aggregation (MLAG) Active-Standby for VPLS feature facilitates implementation of the MLAG Active-Standby between the VPLS PE devices. This facilitates Attachment Circuit (AC) redundancy for the dual-homed Customer Edge (CE) device. This means there is one Active link and another Standby link between a CE device and the PE devices. When an Active link fails, the Standby link becomes Active with the MLAG switch, resulting in a change in topology. The change in topology requires a mac-flush in the peer devices for faster convergence. The feature also facilitates forwarding the automatic mac-flush message to all the peer nodes, reducing the convergence time when a link fails.



The above figure shows a VPLS mesh linked to a dual-homed CE connected to PE1 and PE2. The traffic flows from CE to PE3 through the active MLAG PE1, while PE2 remains a standby MLAG. When PE1 experiences a link failure, the standby MLAG becomes active, and traffic flows through the PE2 device to reach the destination PE3. Configuring the `mac flush send on mlag switchover` command enables MAC flush PDU tx, and PE2

automatically sends the MAC flush message to all the peer nodes PE1, PE3, and PE4 as depicted by the red arrows, which reduces the convergence time.

Consider the following point while configuring this feature:

- Only Active-Standby MLAG is supported.
- Only flush-all-but-mine MAC-flush approach is supported.
- If multiple Attachment Circuits (ACs) are mapped to the same MLAG interface, the MLAG switch is not triggered when only one or a few AC go down.
- MAC flush on a remote VPLS peer is supported only when LDP signaling (not BGP signaling) is used. A MAC flush message (flush-all-but-mine) will have an empty MAC address list. An empty MAC address list will flush all the MAC addresses and cannot flush any selected MAC addresses.

Benefit

This feature enhances the reliability of VPLS by providing redundancy using MLAG.

Prerequisites

Refer to the MLAG Configuration section in the *OcNOS Layer 2 Guide*.


Configuration

This section shows the configuration of the MLAG VPLS with automatic MAC flush capability during the MLAG switchover.

Topology

The following topology shows a VPLS mesh. The traffic is flowing from SW to PE3. PE1 has an MLAG configured as Active and PE2 as Standby. When a link failure occurs between SW and PE1 (the Active MLAG node), the link between SW and PE2 becomes Active.

Figure 85. MLAG VPLS configuration topology



Configuring MLAG on PE1 (Active) and PE2 (Standby)

1. Type `bridge 1 protocol rstp vlan-bridge` to create a bridge and attach it to the RSTP. Type the command `vlan database` to enter the VLAN configuration mode, and then `vlan 100-300 bridge 1`

`state enable` to enable the VLAN 100 to 300 on the bridge 1.

```
(config)#bridge 1 protocol rstp vlan-bridge
(config)#vlan database
(config-vlan)#vlan 100-300 bridge 1 state enable
```

2. Type `mcec domain configuration` to enter the MCEC domain configuration mode and then configure the `domain-address` to identify the MCEC domain. Type the command `domain-system-number 1` to configure the domain system number that uniquely identifies the domain system in MCEC domain and then type the command `intra-domain-peer 2.2.2.2 source-address 9.9.9.9` to map an interface as intra domain peer that connects the domain system with its neighbor in a MCEC domain.

```
(config)#mcec domain configuration
(config-mcec-domain)#domain-address 1111.2222.3333
(config-mcec-domain)#domain-system-number 1
(config-mcec-domain)#intra-domain-peer 2.2.2.2 source-address 9.9.9.9
```

3. Type the command `interface mlag1` and then `switchport` to configure the interface. Type `bridge-group 1` to bind an interface with the bridge group. Type `switchport mode trunk` to set the interface to the trunk port that allows multiple VLAN configurations in the interface, and then type `switchport trunk allowed vlan add 100-150` to configure the required VLAN identifier. Type the command `switchover type revertive 2` to return back to initial MLAG after 2 seconds after fail recovery.

```
(config)#interface mlag1
(config-if)#switchport
(config-if)#bridge-group 1
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan add 100-150
(config-if)#switchover type revertive 2
```

4. Type `interface po1` to enter the specified port channel, type `switchport` to configure the interface as Layer 2, and the type `mlag 1` to enable the specific MLAG.

```
(config)#interface po1
(config-if)#switchport
(config-if)#mlag 1
```

5. Type `interface ce5` to configure the interface and then type `channel-group 1 mode active` to add the interface to the channel group.

```
(config)#interface ce5
(config-if)#channel-group 1 mode active
```

Configuring VPLS Session on PE1 (Active MLAG Peer), PE2 (Standby MLAG Peer), and PE3

1. Type `mpls vpls vpls1 200` to create an VPLS instance. Type `mac-flush-send-on-mlag-switchover` to enable the MAC flush PDU tx during the MLAG switchover, type `signaling ldp` to enter the VPLS signaling mode as only LDP signaling is supported for the MAC flush, and then type `vpls-peer 2.2.2.2` and `vpls-peer 31.31.31.31` to add a peer to a VPLS domain.

```
(config)#mpls vpls vpls1 200
(config-vpls)#mac-flush-send-on-mlag-switchover
(config-vpls)#signaling ldp
(config-vpls-sig)#vpls-peer 2.2.2.2
(config-vpls-sig)#vpls-peer 31.31.31.31
(config-vpls-sig)#exit-signaling
(config-vpls)#exit-vpls
```

2. Type `interface po1.200 switchport` to create a sub-interface. In the sub-interface, type `encapsulation dot1q 200` to select the type of encapsulation as dot1q with the VLAN ID 200. Type

access-if-vpls to create a VPLS access port and then mpls-vpls vpls1 to bind the VPLS instance to the subinterface.

```
(config)#interface pol.200 switchport
(config-if)#encapsulation dot1q 200
(config-if)#access-if-vpls
(config-acc-if-vpls)#mpls-vpls vpls1
```

Running configurations

PE1

```
bridge 1 protocol rstp vlan-bridge
vlan database
vlan 100-300 bridge 1 state enable

mcec domain configuration
domain-address 1111.2222.3333
domain-system-number 1
intra-domain-peer 2.2.2.2 source-address 9.9.9.9

interface mlag1
switchport
bridge-group 1
switchport mode trunk
switchport trunk allowed vlan add 100-150
switchover type revertive 2

interface pol
switchport
mlag 1

mpls vpls vpls1 200
mac-flush-send-on-mlag-switchover
signaling ldp
vpls-peer 2.2.2.2
vpls-peer 31.31.31.31
exit-signaling
exit-vpls

interface pol.200 switchport
encapsulation dot1q 200
access-if-vpls
mpls-vpls vpls1

interface ce5
channel-group 1 mode active
!
```

PE2

```
bridge 1 protocol rstp vlan-bridge

vlan database
vlan 100-300 bridge 1 state enable

mcec domain configuration
domain-address 1111.2222.3333
domain-system-number 1
intra-domain-peer 9.9.9.9 source-address 2.2.2.2
```

```

interface mlag1
  switchport
  bridge-group 1
  switchport mode trunk
  switchport trunk allowed vlan add 100-150
  switchover type revertive 2

interface pol
  switchport
  mlag 1

mpls vpls vpls1 200
  mac-flush-send-on-mlag-switchover
  signaling ldp
  vpls-peer 9.9.9.9
  vpls-peer 31.31.31.31
  exit-signaling
  exit-vpls

interface pol.200 switchport
  encapsulation dot1q 200
  access-if-vpls
  mpls-vpls vpls1

!
interface ce8
  speed 40g
  channel-group 1 mode active
!

```

PE3

```

mpls vpls vpls1 200
  mac-flush-send-on-mlag-switchover
  signaling ldp
  ignore-ac-spoke-state
  vpls-peer 2.2.2.2
  vpls-peer 9.9.9.9
  exit-signaling
  exit-vpls
!
interface xe8.200 switchport
  encapsulation dot1q 200
  access-if-vpls
  mpls-vpls vpls1
!

```

Validation

Use the following show commands to verify the configuration.

Verify MLAG Domain summary on PE1 and PE2 to verify the Current MLAG status and MLAG Synchronization

```

PE1# #show mlag domain summary

-----
Domain Configuration
-----

```

```

Domain System Number      : 1
Domain Address            : 1111.2222.3333
Domain Priority            : 32768
Source Address            : 9.9.9.9
Intra-domain-peer        : 2.2.2.2
Domain Adjacency          : UP
MCEC PDU local version    : 1
MCEC PDU peer version    : 1
Domain Sync via           : Intra-domain-peer

-----
MLAG Configuration
-----

MLAG-1
Mapped Aggregator        : po1
Physical properties Digest : 54 a9 3a 2a 2b 50 65 bb 3c bc 3d bd c2 43 d6 22
Total Bandwidth          : 40g
Mlag Sync                : IN_SYNC
Mode                     : Active-Standby
Current Mlag state       : Active
Switchover-mode          : Revertive (2s)

```

```
PE2#show mlag domain summary
```

```

-----
Domain Configuration
-----

Domain System Number      : 2
Domain Address            : 1111.2222.3333
Domain Priority            : 32768
Source Address            : 2.2.2.2
Intra-domain-peer        : 9.9.9.9
Domain Adjacency          : UP
MCEC PDU local version    : 1
MCEC PDU peer version    : 1
Domain Sync via           : Intra-domain-peer

-----
MLAG Configuration
-----

MLAG-1
Mapped Aggregator        : po1
Physical properties Digest : 54 a9 3a 2a 2b 50 65 bb 3c bc 3d bd c2 43 d6 22
Total Bandwidth          : 40g
Mlag Sync                : IN_SYNC
Mode                     : Active-Standby
Current Mlag state       : Standby
Switchover-mode          : Revertive (2s)
Revert Timer             : OFF

```

Verify VPLS sessions on PE1, PE2, and PE3

```

PE1#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-------------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 200 | 2.2.2.2 | | | 24323 | 25608 | xe10 | 25608 |
| 2/Up | | 1 | LDP | Active | 00:20:29 | - | |
| 200 | 31.31.31.31 | | | 24325 | 25604 | xe10 | 26885 |
| Up | 2 | | LDP | Active | 00:20:29 | - | 2/ |

```
PE2#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-------------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 200 | 9.9.9.9 | | N/A | 25608 | N/A | | 25608 |
| 0/Dn | | 3 | LDP | Standby | - | - | |
| 200 | 31.31.31.31 | | N/A | 25602 | N/A | | 26890 |
| | 4 | | LDP | Standby | - | - | 0/Dn |

```
PE3#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 200 | 2.2.2.2 | | N/A | 26885 | N/A | | 25604 |
| 0/Dn | | 2 | LDP | Standby | - | - | |
| 200 | 9.9.9.9 | | 24320 | 26890 | xe4 | | 25602 |
| 2/Up | | 1 | LDP | Active | 17:35:59 | - | |

Verify VPLS MAC addresses learned on PE1, PE2, and PE3

```
PE1#show mpls vpls mac-address
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | MAC address | Learned from | Vlan-Id | Peer address | Time- |
|---------|----------------|--------------|---------|--------------|-------|
| out | Move | Count | | | |
| 200 | 0000.2000.0020 | xe10 | - | | |
| | 31.31.31.31 | 300 | | | 0 |
| 200 | 0010.2000.0020 | pol.200 | | | |

```
PE2#show mpls vpls mac-address
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | MAC address | Learned from | Vlan-Id | Peer address | Time- |
|---------|----------------|--------------|---------|--------------|-------|
| out | Move | Count | | | |
| 200 | 0000.2000.0020 | xe10 | - | | |
| | 31.31.31.31 | 300 | | | 0 |

```
PE3#show mpls vpls mac-address
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP
```

| VPLS-ID | MAC address | Learned from | Vlan-Id | Peer address | Time- |
|---------|----------------|--------------|---------|--------------|-------|
| out | Move | Count | | | |
| 200 | 0000.2000.0020 | xe8.200 | | | |
| 200 | 0010.2000.0020 | xe5 | - | | |
| | 9.9.9.9 | 296 | | | 0 |

MLAG Active peer went down (Shutdown Active MLAG peer interface on PE1 or on switch)

```
PE1(config)#interface pol1
PE1(config-if)#shutdown
PE1(config-if)#commit
PE1(config)#end
```

Verify MLAG domain summary on PE1 and PE2 to verify the Current MLAG status and MLAG Synchronization

```
PE1#show mlag domain summary
```

```
-----  
Domain Configuration  
-----
```

```
Domain System Number      : 1  
Domain Address            : 1111.2222.3333  
Domain Priority           : 32768  
Source Address            : 9.9.9.9  
Intra-domain-peer        : 2.2.2.2  
Domain Adjacency         : UP  
MCEC PDU local version    : 1  
MCEC PDU peer version     : 1  
Domain Sync via          : Intra-domain-peer
```

```
-----  
MLAG Configuration  
-----
```

```
MLAG-1
```

```
Mapped Aggregator        : po1  
Physical properties Digest : 54 a9 3a 2a 2b 50 65 bb 3c bc 3d bd c2 43 d6 22  
Total Bandwidth          : 40g  
Mlag Sync                : IN_SYNC  
Mode                     : Active-Standby  
Current Mlag state       : Standby  
Switchover-mode          : Revertive (2s)
```

```
PE2#show mlag domain summary
```

```
-----  
Domain Configuration  
-----
```

```
Domain System Number      : 2  
Domain Address            : 1111.2222.3333  
Domain Priority           : 32768  
Source Address            : 2.2.2.2  
Intra-domain-peer        : 9.9.9.9  
Domain Adjacency         : UP  
MCEC PDU local version    : 1  
MCEC PDU peer version     : 1  
Domain Sync via          : Intra-domain-peer
```

```
-----  
MLAG Configuration  
-----
```

```
MLAG-1
```

```
Mapped Aggregator        : po1  
Physical properties Digest : 54 a9 3a 2a 2b 50 65 bb 3c bc 3d bd c2 43 d6 22  
Total Bandwidth          : 40g  
Mlag Sync                : IN_SYNC  
Mode                     : Active-Standby  
Current Mlag state       : Active  
Switchover-mode          : Revertive (10s)  
Revert Timer             : OFF
```

Verify VPLS sessions on PE1, PE2, and PE3

```
PE1#show mpls vpls mesh
```

(m) - Service mapped over multipath transport
 (e) - Service mapped over LDP ECMP

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-------------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 200 | 2.2.2.2 | | N/A | 25608 | N/A | | 25608 0/ |
| Dn | 1 | LDP | Standby | - | - | | |
| 200 | 31.31.31.31 | | N/A | 25604 | N/A | 26885 | 0/Dn |
| | 2 | LDP | Standby | - | - | | |

PE2#show mpls vpls mesh

(m) - Service mapped over multipath transport
 (e) - Service mapped over LDP ECMP

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-------------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 200 | 9.9.9.9 | 24321 | 25608 | ce6 | 25608 | 2/Up | 3 L |
| DP | Active | 00:21:54 | - | | | | |
| 200 | 31.31.31.31 | 24325 | 25602 | ce6 | 26890 | 2/Up | 4 L |
| DP | Active | 17:35:21 | - | | | | |

PE3#show mpls vpls mesh

(m) - Service mapped over multipath transport
 (e) - Service mapped over LDP ECMP

| VPLS-ID | Peer Addr | Tunnel-Label | In-Label | Network-Intf | Out-Label | Lkps/St | PW-INDEX |
|--------------|-----------|--------------|-----------|--------------|-----------|---------|----------|
| SIG-Protocol | Status | UpTime | Ext-Color | | | | |
| 200 | 2.2.2.2 | 24320 | 26890 | xe4 | 25602 | 2/Up | 1 L |
| DP | Active | 17:45:58 | - | | | | |
| 200 | 9.9.9.9 | N/A | 26885 | N/A | 25604 | 0/Dn | 2 L |
| DP | Standby | - | - | | | | |

Verify VPLS MAC addresses learned on PE1, PE2, and PE3

PE1#show mpls vpls mac-address

(m) - Service mapped over multipath transport
 (e) - Service mapped over LDP ECMP

| VPLS-ID | MAC address | Learned from | Vlan-Id | Peer address | Time- |
|---------|----------------|--------------|---------|--------------|-------|
| out | Move | Count | | | |
| 200 | 0000.2000.0020 | xe10 | - | | |
| | 31.31.31.31 | 300 | | | 0 |

PE2#show mpls vpls mac-address

(m) - Service mapped over multipath transport
 (e) - Service mapped over LDP ECMP

| VPLS-ID | MAC address | Learned from | Vlan-Id | Peer address | Time- |
|---------|----------------|--------------|---------|--------------|-------|
| out | Move | Count | | | |
| 200 | 0000.2000.0020 | xe3 | - | | |
| | 31.31.31.31 | 300 | | | 0 |
| 200 | 0010.2000.0020 | po1.200 | | | |

PE3#show mpls vpls mac-address

(m) - Service mapped over multipath transport
 (e) - Service mapped over LDP ECMP

| VPLS-ID | MAC address | Learned from | Vlan-Id | Peer address | Time- |
|---------|----------------|--------------|---------|--------------|-------|
| out | Move | Count | | | |
| 200 | 0000.2000.0020 | xe8.200 | | | |
| 200 | 0010.2000.0020 | xe4 | - | | |
| | 2.2.2.2 | 296 | | | 0 |

MLAG Active-Standby for VPLS commands

The MLAG Active-Standby for VPLS introduces the following configuration command.

switchover type

The revertive time range is revised from <1-255> to <1-3600>. For more details, refer to switchover type command in the Multi-chassis Link Aggregation Commands chapter in the *Layer 2 Guide*.

mac flush send on mlag switchover

Use this command to enable the MAC flush PDU tx during the MLAG switchover.

Use the `no` command to disable the MAC PDU tx flush during the MLAG switchover.

Command Syntax

```
mac-flush-send-on-mlag-switchover
```

```
no mac-flush-send-on-mlag-switchover
```

Parameters

None

Default

Disabled

Command Mode

VPLS mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

This example shows how to configure mac flush automatically on the peer nodes:

```
OcNOS#configure terminal
OcNOS(config)#mpls vpls VPLS-102 102
OcNOS(config-vpls)#mac-flush-send-on-mlag-switchover
OcNOS(config-vpls)#commit
OcNOS(config-vpls)#end
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms/Acronym | Description |
|-------------------|--|
| VPLS | Virtual Private Local Area Network Service is a technology that uses a shared Layer 2 Virtual Private Network (L2VPN) to establish communication between different geographical sites as if they are in the same Local Area Network (LAN). |
| MPLS | Multi-Protocol Label Switching is a high-performance routing protocol that attaches labels to data packets to efficiently forward and route data in a network. |
| CE | Customer Edge is a device that resides at the customer premises and bridges the customer's internal network to the service provider network. |

| Key Terms/Acronym | Description |
|-------------------|--|
| PE | Provider Edge is a device that resides at the provider's edge, manages traffic from the customer edge, and forwards it to the provider's backbone. |
| MLAG | Multi-Chassis Link Aggregation is a network architecture that uses multiple switches or routers to work as a single logical device to increase redundancy and bandwidth. |
| AC | Attachment circuit is an interface or link that connects a CE device to the PE device. |
| MAC | Media Access Control is a unique identifier of a device's network interface card (NIC) to communicate in a local area network (LAN). |

BGP Auto-Discovery (AD) for LDP VPLS

Overview

The BGP Auto-Discovery enables automatic discovery of VPLS peers, eliminating the need for manual peer configuration. Once discovered, pseudo-wires (PWs) between peers are established using LDP signaling, streamlining the VPLS setup process.

OcNOS supports both BGP Auto-Discovery (BGP-AD) and manual VPLS peer configuration within the same VPLS instance. This is beneficial in topologies where certain VPLS peers do not support BGP-AD, as they can still establish VPLS sessions using manual peer configuration.



Note: In BGP AD VPLS, if rd and/or router-target is configured in ASN4:nn format while BGP extended ASN support is disabled or vice-versa then BGP auto-discovery will not work and mesh PW will not come up. In such scenarios user is expected to remove the BGP AD VPLS instance and configure it again with correct rd and route-target config as per the BGP extended ASN support.

Benefits

The BGP Auto-Discovery provides in following aspects:

- Simplifies the VPLS configuration process.
- Enhances network scalability.
- Improves scaling efficiency when used with route reflectors.

Limitations

- MIB is not supported.
- The PW group feature is not supported in OcNOS for both FEC128 and FEC129 signaling in BGP AD VPLS. During interoperability, avoid configuring PW groups on peer devices, as related wildcard label messages will be ignored. However, OcNOS now parses PW Group ID TLVs with a group ID of 0—commonly sent by some vendors—even if PW groups are not configured, ensuring interoperability.
- For using a local IP address for BGP Auto-Discovery (AD) NLRI advertisement and as the SAll in FEC129-based pseudowire (PW) establishment, the system by default selects the secondary IP address on the lo (loopback) interface. If no secondary IP is configured, it falls back to the BGP router ID. However, if a secondary IP is configured on the loopback interface and it differs from the BGP router ID, BGP AD will fail to operate correctly. Additionally, the LDP transport address must match the selected local IP address (i.e., the secondary loopback IP or router ID) to ensure consistency across control-plane protocols.
- Maximum of 32 route-targets are supported per VPLS instance.
- Route-leaking is not supported and for that same route-target cannot be configured in 2 or more VPLS instances.
- Auto-RT and route-map is not supported.
- RD and L2-VPN ID (if configured) must be the same for VPLS peering.

Prerequisites

- **Define Interfaces and Loopback Addresses:**

Configure Layer 2 interfaces, such as port channel interfaces (e.g., po1), and assign IP addresses for identification and routing. Additionally, configure loopback IP addresses to establish key connectivity points. These settings ensure efficient network routing and communication.

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 2.2.2.2/32 secondary  
  ipv6 address ::1/128  
  
interface xe14  
  ip address 30.1.1.2/24
```

- **Configure IGP for Dynamic Routing:**

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

- **ISIS Configuration:**

```
router isis 1  
  is-type level-2-only  
  metric-style wide  
  microloop-avoidance level-2  
  mpls traffic-eng router-id 2.2.2.2  
  mpls traffic-eng level-2  
  capability cspf  
  dynamic-hostname  
  bfd all-interfaces  
  net 49.0000.0000.0002.00  
  passive-interface lo  
  !  
  interface xe14  
    mpls ldp-igp sync isis level-2  
    isis network point-to-point  
    ip router isis 1
```

- **OSPF Configuration:**

```
router ospf 1  
  ospf router-id 2.2.2.2  
  network 2.2.2.2/32 area 0.0.0.0  
  network 30.1.1.0/24 area 0.0.0.0!  
  !  
  interface xe14  
    ip ospf network point-to-point
```

- **LDP Configuration:**

```
router ldp  
  router-id 100.1.1.1  
  transport-address ipv4 100.1.1.1
```

- **BGP Configuration:**

```
!  
router bgp 64000
```

```

bgp router-id 100.1.1.1
neighbor BGP-AD peer-group
neighbor BGP-AD remote-as 64000
neighbor BGP-AD update-source lo
neighbor BGP-AD fall-over bfd multihop
neighbor 150.1.1.1 peer-group BGP-AD
neighbor 160.1.1.1 peer-group BGP-AD
neighbor 170.1.1.1 peer-group BGP-AD
!
address-family l2vpn vpls
neighbor BGP-AD activate
neighbor BGP-AD route-reflector-client
exit-address-family
!
exit
!

```

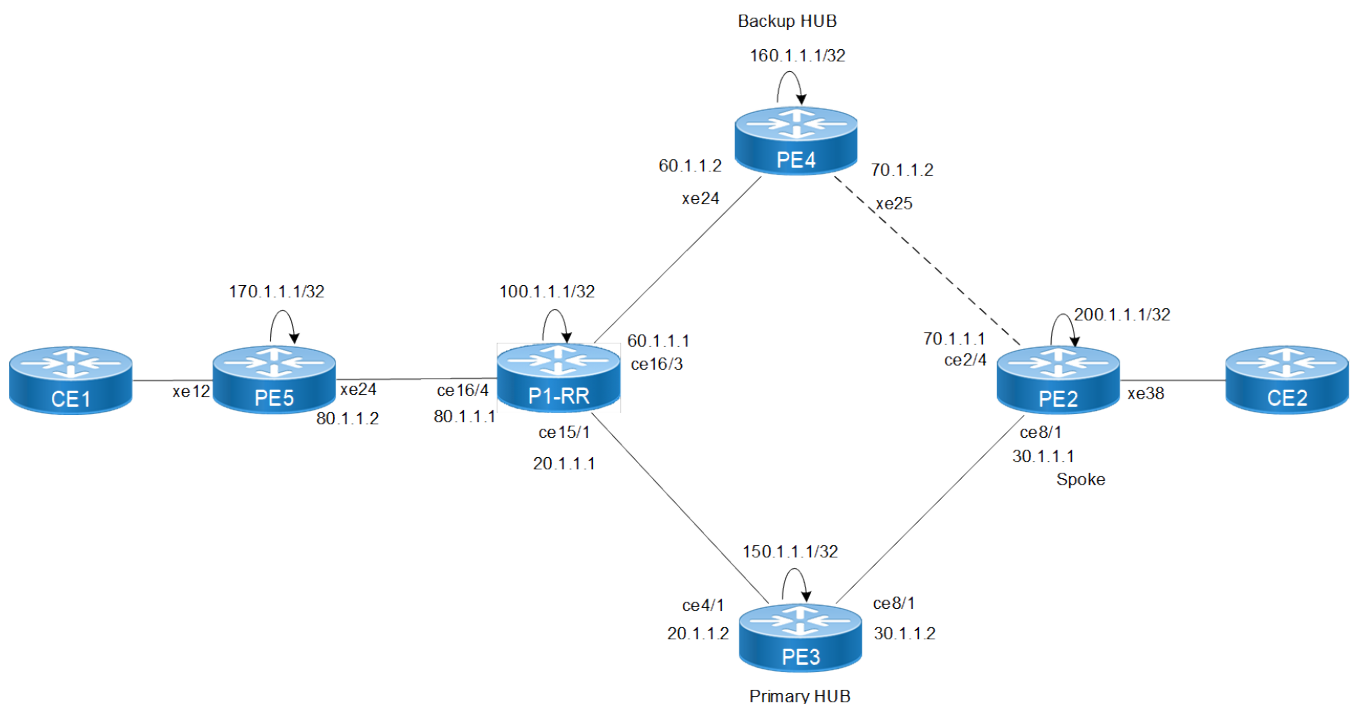
Configuration

Configure various nodes within the topology to set up a H-VPLS session.

Topology

This sample topology provides basic connectivity and routing between the devices.

Figure 86. BGP-A-D LDP VPLS Topology Diagram



PE5

1. Configure router LDP.

```

PE5(config)#router ldp
PE5(config-router)# router-id 170.1.1.1
PE5(config-router)# transport-address ipv4 170.1.1.1
PE5(config-router)# pw-status-tlv

```

2. Configure loopback interface.

```
PE5(config)#interface lo
PE5(config-af)#ip address 170.1.1.1/32 secondary
```

3. Enable LDP and label-switching for core interface.

```
PE5(config)#interface xe24
PE5(config-if)#ip address 80.1.1.2/24
PE5(config-if)#label-switching
PE5(config-if)#mpls ldp-igp sync ospf
PE5(config-if)#ip ospf network point-to-point
PE5(config-if)#enable-ldp ipv4

PE5(config-if)#router ospf 100
PE5(config-if)#network 80.1.1.0/24 area 0.0.0.0
PE5(config-if)#network 170.1.1.1/32 area 0.0.0.0
```

4. Configure BGP.

```
PE5(config)#router bgp 64000
PE5(config-router)#bgp router-id 170.1.1.1
PE5(config-router)#neighbor BGP-AD peer-group
PE5(config-router)#neighbor BGP-AD remote-as 64000
PE5(config-router)#neighbor BGP-AD update-source lo
PE5(config-router)#neighbor BGP-AD fall-over bfd multihop
PE5(config-router)#neighbor 100.1.1.1 peer-group BGP-AD
PE5(config-router)#address-family l2vpn vpls
PE5(config-router-af)#neighbor BGP-AD activate
PE5(config-router-af)#exit-address-family
PE5(config-router)#exit
```

5. Configure VPLS instance.

```
PE5(config)#mpls vpls vpls2 2
PE5(config-vpls)#vpls-mtu 9000
PE5(config-vpls)#signaling ldp
PE5(config-vpls-sig)#vpls-type vlan
PE5(config-vpls-sig)#bgp-auto-discovery
PE5(config-vpls-ldp-sig-bgp-ad)#l2vpn-id 2:1
PE5(config-vpls-ldp-sig-bgp-ad)#rd 2:1
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:1
PE5(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:2
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:3
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:4
PE5(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:5
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:6
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:7
PE5(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:8
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:9
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:10
PE5(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:11
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:12
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:13
PE5(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:14
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:15
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:16
PE5(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:17
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:18
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:19
PE5(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:20
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:21
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:22
PE5(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:23
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:24
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:25
PE5(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:26
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:27
PE5(config-vpls-ldp-sig-bgp-ad)#route-target import 2:28
```

```
PE5(config-vpls-ldp-sig-bgp-ad)# route-target export 10.10.0.2:29
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:30
PE5(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery
PE5(config-vpls-sig)# exit-signaling
PE5(config-vpls)#exit-vpls
```

6. Configure sub-interface and attach vpls-instance to sub-interface.

```
PE5(config)#interface xe12.2 switchport
PE5(config-if)#encapsulation dot1q 2
PE5(config-if)#access-if-vpls
PE5(config-acc-if-vpls)#mpls-vpls vpls2
PE5(config-acc-if-vpls)#exit
PE5(config-if)#exit
```

P1

1. Configure router LDP.

```
P1(config)#router ldp
P1(config-router)#router-id 100.1.1.1
P1(config-router)#transport-address ipv4 100.1.1.1
```

2. Enable LDP and label-switching for core interface.

```
P1(config)#interface ce15/1
P1(config-if)#ip address 20.1.1.1/24
P1(config-if)#label-switching
P1(config-if)#mpls ldp-igp sync ospf
P1(config-if)#ip ospf network point-to-point
P1(config-if)#enable-ldp ipv4
```

3. Configure network interface.

```
P1(config)#interface ce16/3
P1(config-if)#ip address 60.1.1.1/24
P1(config-if)#label-switching
P1(config-if)#mpls ldp-igp sync ospf
P1(config-if)#ip ospf network point-to-point
P1(config-if)#enable-ldp ipv4

P1(config)#interface ce16/4
P1(config-if)#ip address 80.1.1.1/24
P1(config-if)#label-switching
P1(config-if)#mpls ldp-igp sync ospf
P1(config-if)#ip ospf network point-to-point
P1(config-if)#enable-ldp ipv4
```

4. Configure loopback interface.

```
P1(config)#interface lo
P1(config-if)#ip address 100.1.1.1/32 secondary
```

5. Configure OSPF.

```
P1(config)#router ospf 100
P1(config-router)#network 20.1.1.0/24 area 0.0.0.0
P1(config-router)#network 60.1.1.0/24 area 0.0.0.0
P1(config-router)#network 80.1.1.0/24 area 0.0.0.0
P1(config-router)#network 100.1.1.1/32 area 0.0.0.0
```

6. Configure BGP.

```
P1(config)#router bgp 64000
P1(config-router)#bgp router-id 100.1.1.1
P1(config-router)#neighbor BGP-AD peer-group
P1(config-router)#neighbor BGP-AD remote-as 64000
P1(config-router)#neighbor BGP-AD update-source lo
```

```
P1(config-router)#neighbor BGP-AD fall-over bfd multihop
P1(config-router)#neighbor 150.1.1.1 peer-group BGP-AD
P1(config-router)#neighbor 160.1.1.1 peer-group BGP-AD
P1(config-router)#neighbor 170.1.1.1 peer-group BGP-AD
P1(config-router-af)#address-family l2vpn vpls
P1(config-router-af)#neighbor BGP-AD activate
P1(config-router-af)#neighbor BGP-AD route-reflector-client
P1(config-router-af)#exit-address-family
P1(config-router)#exit
```

PE3

1. Configure router LDP.

```
PE3(config)#router ldp
PE3(config-router)# router-id 150.1.1.1
PE3(config-router)# transport-address ipv4 150.1.1.1
```

2. Enable LDP and label-switching for core interface.

```
PE3(config)#interface ce4/1
PE3(config-if)#ip address 20.1.1.2/24
PE3(config-if)#label-switching
PE3(config-if)#mpls ldp-igp sync ospf
PE3(config-if)#ip ospf network point-to-point
PE3(config-if)#enable-ldp ipv4

PE3(config)#interface ce8/1
PE3(config-if)#ip address 30.1.1.2/24
PE3(config-if)#label-switching
PE3(config-if)#mpls ldp-igp sync ospf
PE3(config-if)#ip ospf network point-to-point
PE3(config-if)#enable-ldp ipv4
```

3. Configure loopback interface.

```
PE3(config)#interface lo
PE3(config-if)#ip address 150.1.1.1/32 secondary
```

4. Configure OSPF.

```
PE3(config)#router ospf 100
PE3(config-router)#network 20.1.1.0/24 area 0.0.0.0
PE3(config-router)#network 30.1.1.0/24 area 0.0.0.0
PE3(config-router)#network 150.1.1.1/32 area 0.0.0.0
```

5. Configure BGP.

```
PE3(config)#router bgp 64000
PE3(config-router)#bgp router-id 150.1.1.1
PE3(config-router)#neighbor BGP-AD peer-group
PE3(config-router)#neighbor BGP-AD remote-as 64000
PE3(config-router)#neighbor BGP-AD update-source lo
PE3(config-router)#neighbor BGP-AD fall-over bfd multihop
PE3(config-router)#neighbor 100.1.1.1 peer-group BGP-AD
PE3(config-router-af)#address-family l2vpn vpls
PE3(config-router-af)#neighbor BGP-AD activate
PE3(config-router-af)#neighbor BGP-AD route-reflector-client
PE3(config-router-af)#exit-address-family
PE3(config-router)#exit
```

6. Configure an MPLS L2 Circuit.

```
PE3(config)#mpls l2-circuit vc1 101 200.1.1.1
```

7. Configure an MPLS VPLS Instance.

```
PE3(config)#mpls vpls vpls2 2
```



```

PE3(config-vpls-spoke)#vpls-mtu 9000
PE3(config-vpls-spoke)#vpls-vc vc1
PE3(config-vpls-spoke)#exit-spoke
PE3(config-vpls)#signaling ldp
PE3(config-vpls-sig)#vpls-type vlan
PE3(config-vpls-sig)#bgp-auto-discovery
PE3(config-vpls-ldp-sig-bgp-ad)#l2vpn-id 2:1
PE3(config-vpls-ldp-sig-bgp-ad)#rd 2:1
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:31
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:32
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:33
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:34
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:35
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:36
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:37
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:38
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:39
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:40
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:41
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:42
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:43
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:44
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:45
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:1
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:2
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:3
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:4
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:5
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:6
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:7
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:8
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:9
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:10
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:11
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:12
PE3(config-vpls-ldp-sig-bgp-ad)#route-target import 2:13
PE3(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:14
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:15
PE3(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery
PE3(config-vpls-ldp-sig)#exit-signaling
PE3(config-vpls)#exit-vpls

```

8. Configure the network Interface.

```

PE3(config)#interface ce4/2.2 switchport
PE3(config-if)#encapsulation dot1q 2
PE3(config-if)#access-if-vpls
PE3(config-acc-if-vpls)#mpls-vpls vpls2
PE3(config-acc-if-vpls)#exit
PE3(config-if)#exit

```

PE4

1. Configure router LDP.

```

PE4(config)#router ldp
PE4(config-router)# router-id 160.1.1.1
PE4(config-router)# transport-address ipv4 160.1.1.1

```

2. Configure loopback interface.

```

PE4(config)#interface lo
PE4(config-if)#ip address 160.1.1.1/32 secondary

```

3. Enable LDP and label-switching for core interface.

```

PE4(config)#interface xe24

```

```
PE4(config-if)#ip address 60.1.1.2/24
PE4(config-if)#label-switching
PE4(config-if)#mpls ldp-igp sync ospf
PE4(config-if)#ip ospf network point-to-point
PE4(config-if)#enable-ldp ipv4

PE4(config)#interface xe25
PE4(config-if)#ip address 70.1.1.2/24
PE4(config-if)#label-switching
PE4(config-if)#mpls ldp-igp sync ospf
PE4(config-if)#ip ospf network point-to-point
PE4(config-if)#enable-ldp ipv4
```

4. Configure OSPF.

```
PE4(config)#router ospf 100
PE4(config-router)#network 60.1.1.0/24 area 0.0.0.0
PE4(config-router)#network 70.1.1.0/24 area 0.0.0.0
PE4(config-router)#network 160.1.1.1/32 area 0.0.0.0
```

5. Configure BGP.

```
PE4(config)#router bgp 64000
PE4(config-router)#bgp router-id 160.1.1.1
PE4(config-router)#neighbor BGP-AD peer-group
PE4(config-router)#neighbor BGP-AD remote-as 64000
PE4(config-router)#neighbor BGP-AD update-source lo
PE4(config-router)#neighbor BGP-AD fall-over bfd multihop
PE4(config-router)#neighbor 100.1.1.1 peer-group BGP-AD
PE4(config-router-af)#address-family l2vpn vpls
PE4(config-router-af)#neighbor BGP-AD activate
PE4(config-router-af)#exit-address-family
PE4(config-router)#exit
```

6. Configure an MPLS L2 Circuit.

```
PE4(config)#mpls l2-circuit vc1001 1101 200.1.1.1
```

7. Configure an MPLS VPLS Instance.

```
PE4(config)#mpls vpls vpls2 2
PE4(config-vpls-spoke)#vpls-mtu 9000
PE4(config-vpls-spoke)#vpls-vc vc1001
PE4(config-vpls-spoke)#exit-spoke
PE4(config-vpls-spoke)#signaling ldp
PE4(config-vpls-sig)#vpls-type vlan
PE4(config-vpls-sig)#bgp-auto-discovery
PE4(config-vpls-ldp-sig-bgp-ad)#l2vpn-id 2:1
PE4(config-vpls-ldp-sig-bgp-ad)#rd 2:1
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:1
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:2
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:3
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:4
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:5
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:6
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:7
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:8
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:9
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:10
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:11
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:12
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:13
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:14
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:15
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:16
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:17
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:18
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:19
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:20
```

```

PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:21
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:22
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:23
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:24
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:25
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:26
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:27
PE4(config-vpls-ldp-sig-bgp-ad)#route-target import 2:28
PE4(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:29
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:30
PE4(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery
PE4(config-vpls-ldp-sig)#exit-signaling
PE4(config-vpls)#exit-vpls

```

8. Configure the network Interface.

```

PE4(config)#interface xe12.2 switchport
PE4(config-if)#encapsulation dot1q 2
PE4(config-if)#access-if-vpls
PE4(config-acc-if-vpls)#mpls-vpls vpls2
PE4(config-acc-if-vpls)#exit
PE4(config-if)#exit

```

PE2

1. Configure router LDP.

```

PE4(config)#router ldp
PE4(config-router)# router-id 200.1.1.1
PE4(config-router)# transport-address ipv4 200.1.1.1

```

2. Enable LDP and label-switching for core interface.

```

PE4(config)#interface ce2/4
PE4(config-if)#ip address 70.1.1.2/24
PE4(config-if)#label-switching
PE4(config-if)#mpls ldp-igp sync ospf
PE4(config-if)#ip ospf network point-to-point
PE4(config-if)#enable-ldp ipv4

PE4(config)#interface ce8/1
PE4(config-if)#ip address 30.1.1.2/24
PE4(config-if)#label-switching
PE4(config-if)#mpls ldp-igp sync ospf
PE4(config-if)#ip ospf network point-to-point
PE4(config-if)#enable-ldp ipv4

```

3. Configure loopback interface.

```

PE4(config)#interface lo
PE4(config-if)#ip address 200.1.1.1/32 secondary

```

4. Enable LDP and label-switching for core interface.

```

PE4(config)#interface xe42
PE4(config-if)#ip address 10.1.1.2/24
PE4(config-if)#label-switching
PE4(config-if)#mpls ldp-igp sync ospf
PE4(config-if)#ip ospf network point-to-point
PE4(config-if)#enable-ldp ipv4

```

5. Configure OSPF.

```

PE4(config)#router ospf 100
PE4(config-router)#network 10.1.1.0/24 area 0.0.0.0
PE4(config-router)#network 30.1.1.0/24 area 0.0.0.0
PE4(config-router)#network 70.1.1.1/24 area 0.0.0.0
PE4(config-router)#network 200.1.1.1/32 area 0.0.0.0

```

6. Configure an MPLS L2 Circuit.

```
PE4(config)#mpls l2-circuit vc1 101 150.1.1.1
PE4(config)#mpls l2-circuit vc1001 1101 160.1.1.1
```

7. Configure an MPLS VPLS Instance.

```
PE4(config)#mpls vpls vpls2 2
PE4(config-vpls)#vpls-vc vc1
PE4(config-vpls-spoke)#vpls-vc secondary vc100
PE4(config-vpls-spoke)#exit-spoke
PE4(config-vpls)#exit-vpls
```

8. Configure the network Interface.

```
PE4(config)#interface xe38.2 switchport
PE4(config-if)#encapsulation dot1q 2
PE4(config-if)#access-if-vpls
PE4(config-acc-if-vpls)#mpls-vpls vpls2
PE4(config-acc-if-vpls)#exit
PE4(config-if)#exit
```

Running Configuration**PE5**

```
router ldp
router-id 170.1.1.1
transport-address ipv4 170.1.1.1
!
interface lo
ip address 127.0.0.1/8
ip address 170.1.1.1/32 secondary
ipv6 address ::1/128
!
interface xe24
ip address 80.1.1.2/24
label-switching
mpls ldp-igp sync ospf
ip ospf network point-to-point
enable-ldp ipv4
!
router ospf 100
network 80.1.1.0/24 area 0.0.0.0
network 170.1.1.1/32 area 0.0.0.0
!
router bgp 64000
bgp router-id 170.1.1.1
neighbor BGP-AD peer-group
neighbor BGP-AD remote-as 64000
neighbor BGP-AD update-source lo
neighbor BGP-AD fall-over bfd multihop
neighbor 100.1.1.1 peer-group BGP-AD
!
address-family ipv4 unicast
redistribute connected
neighbor BGP-AD activate
exit-address-family
!
address-family l2vpn vpls
neighbor BGP-AD activate
exit-address-family
!
exit
!
```

```

mpls vpls vpls2 2
vpls-mtu 9000
signaling ldp
vpls-type vlan
bgp-auto-discovery
  l2vpn-id 2:1
  rd 2:1
  route-target import 2:1
  route-target export 10.10.0.2:2
  route-target both 2:3
  route-target import 2:4
  route-target export 10.10.0.2:5
  route-target both 2:6
  route-target import 2:7
  route-target export 10.10.0.2:8
  route-target both 2:9
  route-target import 2:10
  route-target export 10.10.0.2:11
  route-target both 2:12
  route-target import 2:13
  route-target export 10.10.0.2:14
  route-target both 2:15
  route-target import 2:16
  route-target export 10.10.0.2:17
  route-target both 2:18
  route-target import 2:19
  route-target export 10.10.0.2:20
  route-target both 2:21
  route-target import 2:22
  route-target export 10.10.0.2:23
  route-target both 2:24
  route-target import 2:25
  route-target export 10.10.0.2:26
  route-target both 2:27
  route-target import 2:28
  route-target export 10.10.0.2:29
  route-target both 2:30
exit-bgp-auto-discovery
exit-signaling
exit-vpls
!
```

P1

```

router ldp
  router-id 100.1.1.1
  transport-address ipv4 100.1.1.1
!
interface cel5/1
  ip address 20.1.1.1/24
  label-switching
  mpls ldp-igp sync ospf
  ip ospf network point-to-point
  enable-ldp ipv4
!
interface cel6/3
  ip address 60.1.1.1/24
  label-switching
  mpls ldp-igp sync ospf
  ip ospf network point-to-point
  enable-ldp ipv4
!
interface cel6/4
  ip address 80.1.1.1/24
  label-switching
  mpls ldp-igp sync ospf
```

```

ip ospf network point-to-point
enable-ldp ipv4
!
interface lo
ip address 127.0.0.1/8
ip address 100.1.1.1/32 secondary
ipv6 address ::1/128
!
router ospf 100
network 20.1.1.0/24 area 0.0.0.0
network 60.1.1.0/24 area 0.0.0.0
network 80.1.1.0/24 area 0.0.0.0
network 100.1.1.1/32 area 0.0.0.0
!
router bgp 64000
bgp router-id 100.1.1.1
neighbor BGP-AD peer-group
neighbor BGP-AD remote-as 64000
neighbor BGP-AD update-source lo
neighbor BGP-AD fall-over bfd multihop
neighbor 150.1.1.1 peer-group BGP-AD
neighbor 160.1.1.1 peer-group BGP-AD
neighbor 170.1.1.1 peer-group BGP-AD
!
address-family ipv4 unicast
redistribute connected
neighbor BGP-AD activate
exit-address-family
!
address-family l2vpn vpls
neighbor BGP-AD activate
neighbor BGP-AD route-reflector-client
exit-address-family
!
exit
!

```

PE3

```

router ldp
router-id 150.1.1.1
transport-address ipv4 150.1.1.1
!
interface ce4/1
ip address 20.1.1.2/24
label-switching
mpls ldp-igp sync ospf
ip ospf network point-to-point
enable-ldp ipv4
!
interface ce8/1
ip address 30.1.1.2/24
label-switching
mpls ldp-igp sync ospf
ip ospf network point-to-point
enable-ldp ipv4
!
interface lo
ip address 127.0.0.1/8
ip address 150.1.1.1/32 secondary
ipv6 address ::1/128
!
router ospf 100
network 20.1.1.0/24 area 0.0.0.0
network 30.1.1.0/24 area 0.0.0.0
network 150.1.1.1/32 area 0.0.0.0

```

```
!  
router bgp 64000  
  bgp router-id 150.1.1.1  
  neighbor BGP-AD peer-group  
  neighbor BGP-AD remote-as 64000  
  neighbor BGP-AD update-source lo  
  neighbor BGP-AD fall-over bfd multihop  
  neighbor 100.1.1.1 peer-group BGP-AD  
  !  
  address-family ipv4 unicast  
  redistribute connected  
  neighbor BGP-AD activate  
  exit-address-family  
  !  
  address-family l2vpn vpls  
  neighbor BGP-AD activate  
  exit-address-family  
  !  
  exit  
!  
mpls l2-circuit vc1 101 200.1.1.1  
!  
mpls l2-circuit vc1 101 200.1.1.1  
!  
mpls vpls vpls2 2  
  vpls-mtu 9000  
  vpls-vc vc1  
  exit-spoke  
  signaling ldp  
  vpls-type vlan  
  bgp-auto-discovery  
    l2vpn-id 2:1  
    rd 2:1  
    route-target import 2:31  
    route-target export 10.10.0.2:32  
    route-target both 2:33  
    route-target import 2:34  
    route-target export 10.10.0.2:35  
    route-target both 2:36  
    route-target import 2:37  
    route-target export 10.10.0.2:38  
    route-target both 2:39  
    route-target import 2:40  
    route-target export 10.10.0.2:41  
    route-target both 2:42  
    route-target import 2:43  
    route-target export 10.10.0.2:44  
    route-target both 2:45  
    route-target import 2:1  
    route-target export 10.10.0.2:2  
    route-target both 2:3  
    route-target import 2:4  
    route-target export 10.10.0.2:5  
    route-target both 2:6  
    route-target import 2:7  
    route-target export 10.10.0.2:8  
    route-target both 2:9  
    route-target import 2:10  
    route-target export 10.10.0.2:11  
    route-target both 2:12  
    route-target import 2:13  
    route-target export 10.10.0.2:14  
    route-target both 2:15  
  exit-bgp-auto-discovery  
  exit-signaling  
  exit-vpls  
!  
interface ce4/2.2 switchport
```

```
encapsulation dot1q 2
access-if-vpls
  mpls-vpls vpls2
exit
exit
!
```

PE4

```
router ldp
  router-id 160.1.1.1
  transport-address ipv4 160.1.1.1
!
interface lo
  ip address 127.0.0.1/8
  ip address 160.1.1.1/32 secondary
  ipv6 address ::1/128
!
interface xe24
  ip address 60.1.1.2/24
  label-switching
  mpls ldp-igp sync ospf
  ip ospf network point-to-point
  enable-ldp ipv4
!
interface xe25
  ip address 70.1.1.2/24
  label-switching
  mpls ldp-igp sync ospf
  ip ospf network point-to-point
  enable-ldp ipv4
!
router ospf 100
  network 60.1.1.0/24 area 0.0.0.0
  network 70.1.1.0/24 area 0.0.0.0
  network 160.1.1.1/32 area 0.0.0.0
!
router bgp 64000
  bgp router-id 160.1.1.1
  neighbor BGP-AD peer-group
  neighbor BGP-AD remote-as 64000
  neighbor BGP-AD update-source lo
  neighbor BGP-AD fall-over bfd multihop
  neighbor 100.1.1.1 peer-group BGP-AD
!
  address-family ipv4 unicast
  redistribute connected
  neighbor BGP-AD activate
  exit-address-family
!
  address-family l2vpn vpls
  neighbor BGP-AD activate
  exit-address-family
!
  exit
!
mpls l2-circuit vc1001 1101 200.1.1.1
!
mpls vpls vpls2 2
  vpls-mtu 9000
  vpls-vc vc1001
  exit-spoke
  signaling ldp
  vpls-type vlan
  bgp-auto-discovery
```



```

l2vpn-id 2:1
rd 2:1
route-target import 2:1
route-target export 10.10.0.2:2
route-target both 2:3
route-target import 2:4
route-target export 10.10.0.2:5
route-target both 2:6
route-target import 2:7
route-target export 10.10.0.2:8
route-target both 2:9
route-target import 2:10
route-target export 10.10.0.2:11
route-target both 2:12
route-target import 2:13
route-target export 10.10.0.2:14
route-target both 2:15
route-target import 2:16
route-target export 10.10.0.2:17
route-target both 2:18
route-target import 2:19
route-target export 10.10.0.2:20
route-target both 2:21
route-target import 2:22
route-target export 10.10.0.2:23
route-target both 2:24
route-target import 2:25
route-target export 10.10.0.2:26
route-target both 2:27
route-target import 2:28
route-target export 10.10.0.2:29
route-target both 2:30
exit-bgp-auto-discovery
exit-signaling
exit-vpls
!
interface xe12.2 switchport
encapsulation dot1q 2
access-if-vpls
mpls-vpls vpls2
exit
exit
!
end

```

PE2

```

router ldp
router-id 200.1.1.1
transport-address ipv4 200.1.1.1
!
interface ce2/4
ip address 70.1.1.1/24
label-switching
mpls ldp-igp sync ospf
ip ospf network point-to-point
enable-ldp ipv4
!
interface ce8/1
ip address 30.1.1.1/24
label-switching
mpls ldp-igp sync ospf
ip ospf network point-to-point
enable-ldp ipv4
!
interface lo

```

```

ip address 127.0.0.1/8
ip address 200.1.1.1/32 secondary
ipv6 address ::1/128
!
interface xe42
ip address 10.1.1.2/24
label-switching
mpls ldp-igp sync ospf
ip ospf network point-to-point
enable-ldp ipv4
!
router ospf 100
network 10.1.1.0/24 area 0.0.0.0
network 30.1.1.0/24 area 0.0.0.0
network 70.1.1.0/24 area 0.0.0.0
network 200.1.1.1/32 area 0.0.0.0
!
mpls l2-circuit vc1 101 150.1.1.1
!
mpls l2-circuit vc1001 1101 160.1.1.1
!
mpls vpls vpls2 2
vpls-vc vc1
    secondary vc1001
exit-spoke
exit-vpls
!
interface xe38.2 switchport
encapsulation dot1q 2
access-if-vpls
    mpls-vpls vpls2
exit
exit
!

```

Validation

Validate the show output after configuration as shown below.

PE3

```

PE3#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID      Peer Addr      Tunnel-Label  In-Label  Network-Intf  Out-Label  Lkps/St  PW-INDEX
SIG-Protocol  Status        UpTime       Ext-Color
2            160.1.1.1     24321        26240     ce4/1          26240      2/Up     0         L
DP           Active        00:10:29     -
2            170.1.1.1     25602        26241     ce4/1          26240      2/Up     0         L
DP           Active        00:10:16     -
PE3#
PE3#
PE3#show mpls vpls detail
Virtual Private LAN Service Instance: vpls2, ID: 2
SIG-Protocol: LDP
Route-Distinguisher :2:1
Route-Target : export 10.10.0.2:2, export 10.10.0.2:5, export 10.10.0.2:8, import 2:1, both 2:3,
import 2:4, both 2:6, import 2:7, both 2:9,
import 2:10, both 2:12, import 2:40, import 2:31, import 2:13, both 2:42, both 2:33,
both 2:15, import 2:43, import 2:34, both 2:45,
both 2:36, import 2:37, both 2:39, export 10.10.0.2:11, export 10.10.0.2:41, export
10.10.0.2:32, export 10.10.0.2:14,
export 10.10.0.2:44, export 10.10.0.2:35, export 10.10.0.2:38
L2 VPN ID :2:1
Attachment-Circuit: UP

```

```

Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 9000
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: ce4/2.2
Status: Up
Subinterface Match Criteria(s) :
dot1q 2

Mesh Peers:
160.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:10:48)
FEC signaling element: FEC129
FEC129 details:
  agi : 00 0A 00 02 00 00 00 01
  saii: 150.1.1.1
  taii: 160.1.1.1

170.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:10:35)
FEC signaling element: FEC129
FEC129 details:
  agi : 00 0A 00 02 00 00 00 01
  saii: 150.1.1.1
  taii: 170.1.1.1

Spoke Peers:
vc1 (Dn) (Reason: VC on standby)

PE3#
PE3#
PE3#show bgp l2vpn vpls ldp-sig
VPLS-ID      Discovered-Peers  Route-Distinguisher  L2VPN-ID
2            2                2:1                  2:1
PE3#
PE3#
PE3#show bgp l2vpn vpls ldp-sig
VPLS-ID      Discovered-Peers  Route-Distinguisher  L2VPN-ID
2            2                2:1                  2:1
PE3#
PE3#
PE3#show bgp l2vpn vpls ldp-sig detail
VPLS-ID: 2
Local L2VPN-ID      : 2:1
Local RD            : 2:1
Local Route-Targets :
  Import List       : 2:1
                    2:3
                    2:4
                    2:6
                    2:7
                    2:9
                    2:10
                    2:12
                    2:13
                    2:15
                    2:31
                    2:33
                    2:34
                    2:36
                    2:37

```

```

                2:39
                2:40
                2:42
                2:43
                2:45
Export List      : 2:3
                2:6
                2:9
                10.10.0.2:2
                2:12
                10.10.0.2:5
                2:15
                10.10.0.2:8
                10.10.0.2:11
                10.10.0.2:14
                2:33
                2:36
                2:39
                10.10.0.2:32
                2:42
                10.10.0.2:35
                2:45
                10.10.0.2:38
                10.10.0.2:41
                10.10.0.2:44
Discovered Peers : 2
Mesh Peers       :
  BGP Peer-1     : 100.1.1.1
    Peer L2VPN-ID : 2:1
    Peer Route-Targets :
      Export List  : 2:3
                  2:6
                  2:9
                  2:12
                  2:15
                  2:18
                  2:21
                  2:24
                  2:27
                  2:30
                  10.10.0.2:2
                  10.10.0.2:5
                  10.10.0.2:8
                  10.10.0.2:11
                  10.10.0.2:14
                  10.10.0.2:17
                  10.10.0.2:20
                  10.10.0.2:23
                  10.10.0.2:26
                  10.10.0.2:29
    Peer Up time   : 00:11:52
  BGP Peer-2     : 100.1.1.1
    Peer L2VPN-ID : 2:1
    Peer Route-Targets :
      Export List  : 2:3
                  2:6
                  2:9
                  2:12
                  2:15
                  2:18
                  2:21
                  2:24
                  2:27
                  2:30
                  10.10.0.2:2
                  10.10.0.2:5
                  10.10.0.2:8

```

```

10.10.0.2:11
10.10.0.2:14
10.10.0.2:17
10.10.0.2:20
10.10.0.2:23
10.10.0.2:26
10.10.0.2:29
Peer Up time      : 00:12:06

PE3#show mpls vpls spoke
VPLS-ID   Peer Addr   Virtual Circuit  Tunnel-Label  In-Label   Network-Intf   Out-
Label  Lkps/St   Secondary
2      200.1.1.1   vc1              N/A           26242      N/A            26240      2/Dn
---
PE3#show ldp mpls-l2-circuit
Transport  Client      VC      VC      Local      Remote      Destination      Local
Remote
VC ID      Binding    State    Type      VC Label  VC Label  Address           PW
Status
101        VPLS:2     UP       Ethernet  VLAN
26242     26240     200.1.1.1   Forwarding
PE3#

```

PE4

```

PE4#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID   Peer Addr      Tunnel-Label  In-Label   Network-Intf   Out-Label  Lkps/St   PW-INDEX
SIG-Protocol  Status      UpTime      Ext-Color
2           150.1.1.1   24320       26240      xe24            26240      2/Up      0          L
DP          Active      00:12:27    -
2           170.1.1.1   25602       26241      xe24            26241      2/Up      0          L
DP          Active      00:12:14    -
PE4#
PE4#
PE4#
PE4#
PE4#show mpls vpls detail
Virtual Private LAN Service Instance: vpls2, ID: 2
SIG-Protocol: LDP
Route-Distinguisher :2:1
Route-Target : export 10.10.0.2:2, export 10.10.0.2:5, export 10.10.0.2:8, import 2:1, both 2:3,
import 2:4, both 2:6, import 2:7, both 2:9,
import 2:10, both 2:30, both 2:21, both 2:12, import 2:22, import 2:13, both 2:24,
both 2:15, import 2:25, import 2:16, both 2:27,
both 2:18, import 2:28, import 2:19, export 10.10.0.2:20, export 10.10.0.2:11,
export 10.10.0.2:23, export 10.10.0.2:14,
export 10.10.0.2:26, export 10.10.0.2:17, export 10.10.0.2:29
L2 VPN ID :2:1
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 9000
Description: none
service-tpid: dot1.q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe12.2
Status: Up

```

```
Subinterface Match Criteria(s) :
dot1q 2
```

Mesh Peers:

```
150.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:12:35)
FEC signaling element: FEC129
FEC129 details:
  agi : 00 0A 00 02 00 00 00 01
  saii: 160.1.1.1
  taii: 150.1.1.1
```

```
170.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:12:22)
FEC signaling element: FEC129
FEC129 details:
  agi : 00 0A 00 02 00 00 00 01
  saii: 160.1.1.1
  taii: 170.1.1.1
```

Spoke Peers:

```
vc1001 (Up) (UpTime 00:12:17)
```

PE4#

PE4#

PE4#show bgp l2vpn vpls ldp-sig

| VPLS-ID | Discovered-Peers | Route-Distinguisher | L2VPN-ID |
|---------|------------------|---------------------|----------|
| 2 | 2 | 2:1 | 2:1 |

PE4#show bgp l2vpn vpls ldp-sig detail

VPLS-ID: 2

Local L2VPN-ID : 2:1

Local RD : 2:1

Local Route-Targets :

Import List : 2:1

2:3

2:4

2:6

2:7

2:9

2:10

2:12

2:13

2:15

2:16

2:18

2:19

2:21

2:22

2:24

2:25

2:27

2:28

2:30

Export List : 2:3

2:6

2:9

10.10.0.2:2

2:12

10.10.0.2:5

2:15

10.10.0.2:8

2:18

10.10.0.2:11

2:21

10.10.0.2:14

2:24

10.10.0.2:17

2:27

10.10.0.2:20

```

                2:30
                10.10.0.2:23
                10.10.0.2:26
                10.10.0.2:29
Discovered Peers : 2
Mesh Peers      :
  BGP Peer-1    : 100.1.1.1
    Peer L2VPN-ID : 2:1
    Peer Route-Targets :
      Export List : 2:3
                   2:6
                   2:9
                   2:12
                   2:15
                   2:18
                   2:21
                   2:24
                   2:27
                   2:30
                   10.10.0.2:2
                   10.10.0.2:5
                   10.10.0.2:8
                   10.10.0.2:11
                   10.10.0.2:14
                   10.10.0.2:17
                   10.10.0.2:20
                   10.10.0.2:23
                   10.10.0.2:26
                   10.10.0.2:29
    Peer Up time : 00:12:53

```

```

BGP Peer-2      : 100.1.1.1
  Peer L2VPN-ID : 2:1
  Peer Route-Targets :
    Export List : 2:3
                 2:6
                 2:9
                 2:12
                 2:15
                 2:33
                 2:36
                 2:39
                 2:42
                 2:45
                 10.10.0.2:2
                 10.10.0.2:5
                 10.10.0.2:8
                 10.10.0.2:11
                 10.10.0.2:14
                 10.10.0.2:32
                 10.10.0.2:35
                 10.10.0.2:38
                 10.10.0.2:41
                 10.10.0.2:44
    Peer Up time : 00:13:05

```

PE4#

PE4#show mpls vpls spoke

| VPLS-ID | Peer Addr | Virtual Circuit | Tunnel-Label | In-Label | Network-Intf | Out- |
|---------------|-----------|-----------------|--------------|----------|--------------|------------|
| Label Lkps/St | Secondary | | | | | |
| 2 | 200.1.1.1 | vc1001 | 0 | 26242 | xe25 | 26241 2/Up |

PE4#show ldp mpls-l2-circuit

| Transport | Client | VC | VC | Local | Remote | Destination | Local |
|-----------|---------|-----------|---------------|----------|----------|-------------|-------|
| VC ID | Binding | State | Type | VC Label | VC Label | Address | PW |
| Status | | PW Status | | | | | |
| 1101 | VPLS:2 | UP | Ethernet VLAN | | | | |

| 26242 | 26241 | 200.1.1.1 | Forwarding | Forwarding |
|-------|-------|-----------|------------|------------|
| PE4# | | | | |

PE5

```

PE5#show mpls vpls mesh
(m) - Service mapped over multipath transport
(e) - Service mapped over LDP ECMP

VPLS-ID      Peer Addr      Tunnel-Label  In-Label  Network-Intf  Out-Label  Lkps/St  PW-INDEX
SIG-Protocol  Status      UpTime      Ext-Color
2            150.1.1.1    25600       26240     xe24          26241      2/Up     0         L
DP           Active      00:13:07    -
2            160.1.1.1    25601       26241     xe24          26241      2/Up     0         L
DP           Active      00:13:07    -
PE5#
PE5#
PE5#show mpls vpls detail
Virtual Private LAN Service Instance: vpls2, ID: 2
  SIG-Protocol: LDP
    Route-Distinguisher :2:1
    Route-Target : export 10.10.0.2:2, export 10.10.0.2:5, export 10.10.0.2:8, import 2:1, both 2:3,
import 2:4, both 2:6, import 2:7, both 2:9,
import 2:10, both 2:30, both 2:21, both 2:12, import 2:22, import 2:13, both 2:24,
both 2:15, import 2:25, import 2:16, both 2:27,
both 2:18, import 2:28, import 2:19, export 10.10.0.2:20, export 10.10.0.2:11,
export 10.10.0.2:23, export 10.10.0.2:14,
export 10.10.0.2:26, export 10.10.0.2:17, export 10.10.0.2:29
  L2 VPN ID :2:1
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, VPLS Type: Ethernet VLAN, Configured MTU: 9000
Description: none
service-tpid: dot1q
Operating mode: Tagged
Svlan Id: 0
Svlan Tpid: 8100
MAC Withdrawal:

Configured interfaces:
Interface: xe12.2
Status: Up
Subinterface Match Criteria(s) :
dot1q 2

Mesh Peers:
150.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:13:20)
FEC signaling element: FEC129
FEC129 details:
  agi : 00 0A 00 02 00 00 00 01
  saii: 170.1.1.1
  taii: 150.1.1.1

160.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: None) (Up) (UpTime: 00:13:20)
FEC signaling element: FEC129
FEC129 details:
  agi : 00 0A 00 02 00 00 00 01
  saii: 170.1.1.1
  taii: 160.1.1.1

PE5#
PE5#
PE5#show bgp l2vpn vpls ldp-sig
VPLS-ID      Discovered-Peers  Route-Distinguisher  L2VPN-ID

```



```

2                2                2:1                2:1
PE5#
PE5#
PE5#show bgp l2vpn vpls ldp-sig detail
VPLS-ID: 2
Local L2VPN-ID      : 2:1
Local RD            : 2:1
Local Route-Targets :
    Import List     : 2:1
                     2:3
                     2:4
                     2:6
                     2:7
                     2:9
                     2:10
                     2:12
                     2:13
                     2:15
                     2:16
                     2:18
                     2:19
                     2:21
                     2:22
                     2:24
                     2:25
                     2:27
                     2:28
                     2:30
    Export List      : 2:3
                     2:6
                     2:9
                     10.10.0.2:2
                     2:12
                     10.10.0.2:5
                     2:15
                     10.10.0.2:8
                     2:18
                     10.10.0.2:11
                     2:21
                     10.10.0.2:14
                     2:24
                     10.10.0.2:17
                     2:27
                     10.10.0.2:20
                     2:30
                     10.10.0.2:23
                     10.10.0.2:26
                     10.10.0.2:29
Discovered Peers    : 2
Mesh Peers         :
    BGP Peer-1      : 100.1.1.1
        Peer L2VPN-ID : 2:1
        Peer Route-Targets :
            Export List : 2:3
                         2:6
                         2:9
                         2:12
                         2:15
                         2:18
                         2:21
                         2:24
                         2:27
                         2:30
                         10.10.0.2:2
                         10.10.0.2:5
                         10.10.0.2:8
                         10.10.0.2:11
                         10.10.0.2:14

```

```

10.10.0.2:17
10.10.0.2:20
10.10.0.2:23
10.10.0.2:26
10.10.0.2:29
Peer Up time      : 00:13:55

BGP Peer-2       : 100.1.1.1
Peer L2VPN-ID    : 2:1
Peer Route-Targets :
  Export List    : 2:3
                  2:6
                  2:9
                  2:12
                  2:15
                  2:33
                  2:36
                  2:39
                  2:42
                  2:45
                  10.10.0.2:2
                  10.10.0.2:5
                  10.10.0.2:8
                  10.10.0.2:11
                  10.10.0.2:14
                  10.10.0.2:32
                  10.10.0.2:35
                  10.10.0.2:38
                  10.10.0.2:41
                  10.10.0.2:44
Peer Up time      : 00:13:55

```

PE5#

PE2

```

PE2#show mpls vpls detail
Virtual Private LAN Service Instance: vpls2, ID: 2
SIG-Protocol: N/A
Attachment-Circuit: UP
Learning: Enabled
Control-Word: Disabled
Flow Label Status: Disabled, Direction: None, Static: No
Group ID: 0, Configured MTU: 9000
Description: none
service-tpid: dot1q
Operating mode: Raw
MAC Withdrawal:

```

```

Configured interfaces:
Interface: xe38.2
Status: Up
Subinterface Match Criteria(s) :
dot1q 2

```

```

Spoke Peers:
vc1 (Dn) (Reason: VC on standby)
Secondary: vc1001 (Up) (UpTime 00:14:04)

```

PE2#

```

PE2#show mpls vpls spoke
VPLS-ID   Peer Addr   Virtual Circuit  Tunnel-Label  In-Label  Network-Intf  Out-
Label     Lkps/St    Secondary
2         150.1.1.1   vc1              N/A           26240     N/A           26242       0/Dn
vc1001

```

```
2      160.1.1.1    vc1001      0      26241      ce2/4      26242      2/Up
---
PE2#
PE2#
PE2#
PE2#
PE2#show ldp mpls-l2-circuit
Transport      Client      VC      VC      Local      Remote      Destination      Local
Remote
VC ID      Binding      State      Type      VC Label      VC Label      Address      PW
Status
1101      VPLS:2      UP      Ethernet VLAN
26241      26242      160.1.1.1      Forwarding
101      VPLS:2      STBY      Ethernet VLAN
26240      26242      150.1.1.1      Stby      Forwarding
PE2#
```

CLI Commands

The BGP Auto-Discovery (AD) introduces the following configuration commands.

bgp-auto-discovery

Use this command to enable BGP Auto-Discovery for LDP peers.

Use `no` parameter of this command to disable BGP Auto-Discovery for LDP peers.

Command Syntax

```
bgp-auto-discovery
```

```
no bgp-auto-discovery
```

Parameters

None

Default

Disabled

Command Mode

SIGNALING LDP mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Explain or describe the example.

```
#configure terminal
(config)#mpls vpls VPLS100 100
(config-vpls)#signaling ldp
(config-vpls-sig)#bgp-auto-discovery
(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery
(config-vpls-sig)#exit
(config-vpls)#exit
```

l2vpn-id

Use this command to assign a Layer 2 VPN ID for the LDP VPLS.

Use `no` form command to remove the Layer 2 VPN ID configuration.

Command Syntax

```
l2vpn-id ASN:nn_or_IP-address:nn  
no l2vpn-id ASN:nn_or_IP-address:nn
```

Parameters

ASN:nn_or_IP-address:nn

AS number and an arbitrary number (for example, 100:1). Otherwise, specify a 32-bit IP address and an arbitrary number (for example, 192.16.10.1:1).

Default

The default value is set to ASN:VPLS ID when BGP external ASN support is disabled and the BGP ASN is less than 65535; otherwise, specifying a value is mandatory.

Command Mode

BGP AUTO DISCOVERY mode

Applicability

This command is introduced in OcNOS version 6.6.0.

Examples

```
#configure terminal  
(config)#mpls vpls VPLS100 100  
(config-vpls)#signaling ldp  
(config-vpls-sig)#bgp-auto-discovery  
(config-vpls-ldp-sig-bgp-ad)#l2vpn-id 3.3.3.3:3333  
(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery  
(config-vpls-sig)#exit  
(config-vpls)#exit
```

rd (route distinguisher)

Use this command to assign a route distinguisher (RD) for the BGP AD VPLS. The route distinguisher value must be unique within all BGP AD VPLS instances on the router.



Note: BGP auto-discovery requires an RD configuration. Once configured, the RD can be modified but not removed individually; you must remove the entire bgp-auto-discovery configuration block.

Command Syntax

```
rd ASN:nn_or_IP-address:nn
```

Parameters

ASN:nn_or_IP-address:nn

AS number and an arbitrary number (for example, 100:1). Otherwise, specify a 32-bit IP address and an arbitrary number (for example, 192.16.10.1:1).

Default

None

Command Mode

BGP AUTO DISCOVERY mode

Applicability

This command is introduced in OcNOS version 6.6.0.

Examples

```
#configure terminal
(config)#mpls vpls VPLS100 100
(config-vpls)#signaling ldp
(config-vpls-sig)#bgp-auto-discovery
(config-vpls-ldp-sig-bgp-ad)#rd 1.1.1.1:1111
(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery
(config-vpls-sig)#exit
(config-vpls)#exit
```

route-target

Use this command to configure a route-target of type import, export, or both to the BGP AD VPLS.

Command Syntax

```
route-target (import|export|both) (ASN:nn_or_IP-address:nn)
```

Parameters

import

Import routing information.

export

Export routing information.

both

Import and export routing information.

ASN:nn_or_IP-address:nn

AS number and an arbitrary number (for example, 100:1). Otherwise, specify a 32-bit IP address and an arbitrary number (for example, 192.16.10.1:1).

Default

None

Command Mode

BGP AUTO DISCOVERY mode

Applicability

This command is introduced in OcNOS version 6.6.0.

Examples

The following example for configuring BGP Auto-Discovery, including route-target import/export:

```
#configure terminal
(config)#mpls vpls vpls2 2
(config-vpls)#vpls-mtu 9000
(config-vpls)#signaling ldp
(config-vpls-sig)#vpls-type vlan
(config-vpls-sig)#bgp-auto-discovery
(config-vpls-ldp-sig-bgp-ad)#l2vpn-id 2:1
(config-vpls-ldp-sig-bgp-ad)#rd 2:1
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:1
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:2
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:3
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:4
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:5
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:6
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:7
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:8
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:9
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:10
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:11
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:12
```



```
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:13
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:14
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:15
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:16
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:17
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:18
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:19
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:20
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:21
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:22
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:23
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:24
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:25
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:26
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:27
(config-vpls-ldp-sig-bgp-ad)#route-target import 2:28
(config-vpls-ldp-sig-bgp-ad)#route-target export 10.10.0.2:29
(config-vpls-ldp-sig-bgp-ad)#route-target both 2:30
(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery
(config-vpls-sig)#exit-signaling
(config-vpls)#exit-vpls
```

Ethernet Data Plane Loopback (EDPL)

Overview

Ethernet Data Plane Loopback (EDPL) enables remote testing of Ethernet interfaces using the RFC 2544 benchmarking methodology. In this test, a host sends packets to a reflector, which loops them back after swapping specific fields. This process helps measure key performance metrics such as throughput, latency, frame loss, and burst handling of network devices.

Feature Characteristics

- **RFC 2544 Benchmarking:** EDPL follows RFC 2544 guidelines to evaluate device performance and verify compliance with Service Level Agreements (SLAs).
- Includes Throughput, Back-to-Back, Frame Loss, and Latency measurements to assess forwarding, buffering, and delay performance.
- The host acts as a reflector, looping back test packets from a host.
 - L2 traffic: Swaps source and destination MAC addresses.
 - L3 traffic: Swaps IP addresses, MAC addresses, and/or UDP/TCP ports (only in upstream direction).
- Supports upstream (egress) and downstream (ingress) loopbacks on physical ports, sub-interfaces, LAGs, LAG sub-interfaces, and logical interfaces (IRB, VLAN, BVI). Logical interfaces require an associated L2 port.
- EDPL can be configured on physical ports, sub-interfaces, LAGs, LAG sub-interfaces, and logical interfaces (IRB, VLAN, BVI). Logical interfaces require defining the corresponding L2 port.

Benefits

- **Simplified Performance Testing:** Enables quick and remote validation of Ethernet interface performance without requiring external packet generators at every test point.
- **Standards-Based Verification:** Uses RFC 2544 methodology, ensuring consistent and industry-recognized benchmarking across different vendors and devices.
- **Comprehensive Metrics:** Provides detailed insights into throughput, latency, frame loss, and buffer performance, helping operators evaluate real network behavior.
- **Operational Efficiency:** Reduces downtime and testing complexity by allowing on-device reflector configuration through simple CLI commands.
- **Flexible Deployment:** Supports multiple loopback directions and interface types (physical, logical, and aggregated), making it adaptable to various network architectures and test scenarios.

Prerequisites

- **Device Operational:** The router or switch must be powered on and fully operational.
- **Routing Functionality:** All routing protocols and configurations should be active and functioning correctly.
- **End-to-End Connectivity:** Data traffic must flow successfully between source and destination nodes before initiating the EDPL test.
- **Interface Readiness:** The interfaces intended for EDPL testing should be in an up/up state and properly configured.

- For upstream EDPL tests, enable appropriate filter groups based on the address family.
- Traffic Type Validation: Confirm that the host uses unicast packets only, as EDPL does not loop back multicast or broadcast frames.

Limitation

- Platform Support: The feature is supported only on Jericho2 and J2C+ devices.
- Reflector Function Only: The host function is not supported in this release — only the reflector role is implemented.
- Traffic Restriction: Set the unicast packet's destination MAC address as destination-mac in the EDPL profile to prevent loopback multicast and broadcast traffic.
- Directional Limitation: L3 reflector operation is supported only in the upstream direction (egress side).
- EDPL for VxLAN upstream, IPv6 loopback is not supported by Q2.
- Downstream EDPL packet counter is available only on the physical port, even when traffic goes to Sub-interface.


Configuration

Configure the Ethernet Data Plane Loopback (EDPL) feature by creating an EDPL profile that defines the loopback parameters such as direction, address family, and actions, and an EDPL test that specifies the profile, test duration, and source interface for the loopback operation. You can configure an EDPL profile for IPv4, IPv6, or L2 traffic.

Topology

The topology represents a service provider network for Ethernet Data Plane Loopback (EDPL) testing, connecting Host-1 and Host-2 through customer edge (CE1, CE2), provider edge (PE1, PE2), and a core (P) router. Traffic flows from Host-1 to Host-2 across CE–PE–P–PE–CE, with PE routers acting as EDPL reflectors to loop back test packets. This setup supports L2, IPv4, and IPv6 loopback validation, enabling measurement of throughput, latency, and frame loss to verify end-to-end data plane performance.

Topology 1. EDPL Configuration



Enable EDPL for IPv4

The following configuration enables Ethernet Data Plane Loopback (EDPL) for IPv4 traffic in the upstream direction on the PE2, allowing reflector functionality to loop back IPv4 packets as per RFC 2544.

1. Create an EDPL profile for IPv4: Define the loopback direction and address family. Enable UDP/TCP port swapping to allow packet headers to be modified and looped back to the host.

```
(config)# edpl profile edpl-1
(edpl-profile-config)# direction upstream
(edpl-profile-config)# address-family ipv4
(edpl-profile-config)# udp-tcp-port-swap
(edpl-profile-config)# commit
(edpl-profile-config)# end
```

2. (Optional) Configure IPv4 filters: Configure specific match criteria to loop back only selected IPv4 packets based on IP address and UDP ports.

```
(config)# edpl profile edpl-1
(edpl-profile-config)# source-ip-address 10.1.1.1
(edpl-profile-config)# destination-ip-address 50.1.1.1
(edpl-profile-config)# source-udp-port 1000
(edpl-profile-config)# destination-udp-port 5000
(edpl-profile-config)# commit
(edpl-profile-config)# end
```

3. Create and bind the EDPL test: Associate the IPv4 profile with a test, define the duration, and specify the source interface used for loopback.

```
(config)# edpl test test-ipv4
(edpl-test)# edpl-profile edpl-1
(edpl-test)# duration hours 1
(edpl-test)# source-port ce2
(edpl-test)# commit
(edpl-test)# end
```

4. Start and stop the IPv4 test

```
(exec)# edpl start test-ipv4
(exec)# edpl stop test-ipv4
```

Enable EDPL for IPv6

The following configuration enables EDPL for IPv6 traffic, allowing upstream reflector functionality to loop back IPv6 packets with IPv6 addresses and UDP/TCP port swapping.

1. Create an EDPL profile for IPv6: Define the loopback direction and address family. Enable UDP/TCP port swapping to allow packet headers to be modified and looped back to the host.

```
(config)# edpl profile edpl-1
(edpl-profile-config)# direction upstream
(edpl-profile-config)# address-family ipv6
(edpl-profile-config)# udp-tcp-port-swap
(edpl-profile-config)# commit
(edpl-profile-config)# end
```

2. (Optional) Configure IPv6 filters: Configure specific match criteria to loop back only selected IPv4 packets based on IP address and UDP ports.

```
(config)# edpl profile edpl-1
(edpl-profile-config)# source-ip-address 1000::1
(edpl-profile-config)# destination-ip-address 5000::1
(edpl-profile-config)# source-udp-port 1000
(edpl-profile-config)# destination-udp-port 5000
(edpl-profile-config)# commit
(edpl-profile-config)# end
```

3. Create and bind the EDPL test: Associate the IPv6 profile with a test, define the duration, and specify the source interface used for loopback.

```
(config)# edpl test test-ipv6
(edpl-test)# edpl-profile edpl-1
(edpl-test)# duration minutes 30
(edpl-test)# source-port ce2
(edpl-test)# commit
(edpl-test)# end
```

4. Start and stop the IPv6 test.

```
(exec)# edpl start test-ipv6
(exec)# edpl stop test-ipv6
```

Enable EDPL for L2

The following configuration enables EDPL for Layer 2 traffic, allowing the reflector to loop back Ethernet frames by swapping source and destination MAC addresses.

1. Create an EDPL profile for L2: Define the loopback direction and address family.

```
(config)#edpl profile edpl-1
(edpl-profile-config)#direction upstream
(edpl-profile-config)#address-family l2
(edpl-profile-config)#commit
```

2. (Optional) Configure L2 filters: Limit loopback to packets matching specific L2 parameters such as MAC address, VLAN ID, and ethertype.

```
(config)# edpl profile edpl-1
(edpl-profile-config)# destination-mac aa00.bbbb.cccc
(edpl-profile-config)# vlan 10
(edpl-profile-config)# ethertype 0x800
(edpl-profile-config)# commit
(edpl-profile-config)# end
```

3. Create and bind the EDPL test: Associate the L2 profile with a test, define the duration, and specify the interface where loopback occurs.

```
(config)# edpl test test-12
(edpl-test)# edpl-profile edpl-1
(edpl-test)# duration minutes 10
(edpl-test)# source-port ce2
(edpl-test)# commit
(edpl-test)# end
```

4. Start and stop the L2 test.

```
(exec)# edpl start test-12
(exec)# edpl stop test-12
```

Validation

Use the following commands to verify that the Ethernet Data Plane Loopback (EDPL) feature is configured and operating correctly. Confirm that the reflector is active, the correct profile and direction are applied, and the loopback is functioning as expected.

1. Verify EDPL Test Status: Check the active test, loopback direction, and remaining duration.

```
DUT3# show ethernet edpl
Name                Direction    Duration
-----
test-1              Upstream    01:00:00 (00:59:57 remaining)
DUT3#
```

2. View Detailed EDPL Test Information: Display detailed configuration, including applied filters, interface, and swap actions.

```
DUT3# show ethernet edpl detail
test-1 (ce2)
-----
Direction          Upstream
Filters             ipv4-family
                   ce2
Action              ip-swap, udp-tcp-port-swap
Duration            01:00:00 (00:59:56 remaining)
DUT3#
```

3. (Optional) Validate Filter Configuration: Confirm that the configured filters are correctly applied for IPv4, IPv6, or L2 traffic, ensuring that only specified packets are looped back.

a. IPv4 Filter Example

```
DUT3(config)# edpl profile edpl-1
DUT3(edpl-profile-config)# direction upstream
DUT3(edpl-profile-config)# address-family ipv4
DUT3(edpl-profile-config)# udp-tcp-port-swap
DUT3(edpl-profile-config)# source-ip-address 10.1.1.1
DUT3(edpl-profile-config)# destination-ip-address 50.1.1.1
DUT3(edpl-profile-config)# source-udp-port 1000
DUT3(edpl-profile-config)# destination-udp-port 5000
DUT3(edpl-profile-config)# commit
DUT3(edpl-profile-config)# end
```

b. IPv6 Filter Example

```
DUT3(config)# edpl profile edpl-1
DUT3(edpl-profile-config)# direction upstream
DUT3(edpl-profile-config)# address-family ipv6
DUT3(edpl-profile-config)# udp-tcp-port-swap
DUT3(edpl-profile-config)# source-ip-address 1000::1
DUT3(edpl-profile-config)# destination-ip-address 5000::1
DUT3(edpl-profile-config)# source-udp-port 1000
DUT3(edpl-profile-config)# destination-udp-port 5000
DUT3(edpl-profile-config)# commit
DUT3(edpl-profile-config)# end
```

c. L2 Filter Example

```
DUT3(config)# edpl profile edpl-1
DUT3(edpl-profile-config)# direction upstream
DUT3(edpl-profile-config)# address-family l2
DUT3(edpl-profile-config)# destination-mac aa00.bbbb.cccc
DUT3(edpl-profile-config)# vlan 10
DUT3(edpl-profile-config)# ethertype 0x800
DUT3(edpl-profile-config)# commit
DUT3(edpl-profile-config)# end
```

4. Confirm the test behavior on host.

- a. Verify that all transmitted packets from the host are received back without loss.
- b. Ensure IP/MAC/UDP/TCP swaps occur as configured in the EDPL profile.
- c. Check packet counts and latency results to confirm proper reflector functionality.

5. Ensure RX and TX counters increment as expected during test operation.

```
DUT3# show interface ce2 counters
```

Implementation Examples

Here are example scenarios and use cases for implementing Ethernet Data Plane Loopback (EDPL) in a network to validate performance, loopback traffic, and ensure end-to-end data path integrity.

Scenario 1: Validating Data Plane Performance for IPv4 Traffic

In a large-scale service provider network, validating IPv4 data plane performance is crucial before enabling customer services. The operator needs to verify throughput, latency, and frame loss on core links between two routers without impacting live traffic.

Use Case 1: Implementing EDPL for IPv4 upstream loopback allows the operator to perform RFC 2544-compliant testing by using an external host. The host sends IPv4 packets to the PE2 configured as an EDPL reflector. The PE2 loops back the packets after swapping Mac addresses, IP addresses, and/or UDP/TCP ports, allowing the host to measure performance metrics accurately. This approach helps ensure that the network forwarding path, QoS policies, and hardware buffers operate efficiently before service activation.

Scenario 2: IPv6 Transport Verification Across MPLS Core

An operator deploying IPv6 transport services over MPLS needs to validate the path integrity and latency for IPv6 data forwarding between aggregation and core routers. Manual testing is complex and time-consuming.

Use Case 2: Implementing EDPL for IPv6 upstream loopback allows IPv6 packets to be reflected by the PE2 with source and destination IP address swapping. The test verifies correct IPv6 forwarding behavior, confirms label operations, and measures packet loss and round-trip delay. This ensures IPv6 MPLS transport services meet SLA benchmarks and validates that IPv6 forwarding works seamlessly across provider edge (PE) and core devices.

Scenario 3: Ethernet Frame Validation on Access Links (L2)

A data center network operator must verify L2 connectivity and frame forwarding between TOR (Top-of-Rack) switches before enabling VLAN-based tenant services.

Use Case 3: Implementing EDPL for L2 downstream loopback allows the PE2 to reflect Ethernet frames by swapping source and destination MAC addresses. The test ensures the VLAN is properly configured, confirms L2 reachability, and detects potential misconfigurations such as VLAN mismatches or frame tagging errors. This setup is particularly useful during pre-deployment validation or post-maintenance link verification in data centers or access aggregation networks.

Scenario 4: Logical Interface (SVI/IRB/BVI) Validation in Multi-VRF Networks

In multi-VRF environments, where routing and bridging coexist (IRB/BVI interfaces), engineers must confirm traffic forwarding between logical interfaces and their parent physical interfaces.

Use Case 4: Implementing EDPL with logical interfaces (SVI/IRB/BVI) allows engineers to map the logical interface to a parent L2 port using the l2-port option. This ensures that packets are properly looped back at the correct interface layer. By testing through logical interfaces, the operator can validate bridged and routed traffic paths, verify VRF separation, and confirm that data plane forwarding aligns with configured routing policies.

Scenario 5: Pre-Service SLA Verification for New Circuits

Before activating new enterprise or wholesale connections, the operator must confirm that the circuit meets SLA parameters such as latency and frame loss.

Use Case 5: By implementing EDPL reflector tests using IPv4 or L2 upstream configurations, the operator can perform SLA validation without external probes at each end. This reduces operational cost, automates test cycles, and confirms that each link segment meets RFC 2544 benchmarks before provisioning customer traffic.

EDPL CLI Commands

The EDPL introduces the following new configuration commands in OcNOS.

address-family

Use this command to define the address family for the EDPL test. The EDPL feature supports L2, IPv4, and IPv6 address families.

Use the *no* form of this command to remove the configured address family.

Command Syntax

```
address-family <l2 | ipv4 | ipv6>  
no address-family
```

Parameters

l2

Configures Layer 2 loopback.

ipv4

Configures IPv4 loopback.

ipv6

Configures IPv6 loopback.

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

IPv4

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# address-family ipv4  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```

IPv6

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# address-family ipv6  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```

L2

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# address-family l2  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```

direction

Use this command to configure the direction of the EDPL test, determining where loopback occurs within the forwarding path.

Use the *no* form of this command to reset the direction to the default value.

Command Syntax

```
direction <upstream | downstream>
no direction
```

Parameters

upstream

Loops packets at the egress interface.

downstream

Loops packets at the ingress interface.

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1
OcNOS(edpl-profile-config)# direction upstream
OcNOS(edpl-profile-config)# commit
OcNOS(edpl-profile-config)# end
OcNOS#
```

destination-mac

Use this command to specify the destination MAC address for L2 loopback traffic filtering.

Use the *no* form of this command to remove the configured destination MAC address.

Command Syntax

```
destination-mac <mac-address>  
no destination-mac
```

Parameters

<mac-address>

Destination MAC address in the format xxxx.xxxx.xxxx.

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# destination-mac aa00.bbbb.cccc  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```

destination-udp-port

Use this command to configure the destination UDP port number used for filtering packets in EDPL tests.

Use the *no* form of this command to remove the configured destination UDP port.

Command Syntax

```
destination-udp-port <port-number>  
no destination-udp-port
```

Parameters

<port-number>

Source UDP port number (1–65535).

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# destination-udp-port 5000  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```

destination-ip-address

Use this command to specify the destination IP address for filtering IPv4 or IPv6 packets during EDPL tests.

Use the *no* form of this command to remove the configured destination IP filter.

Command Syntax

```
destination-ip-address <ip-address>  
no destination-ip-address
```

Parameters

<ip-address>

Source IPv4 or IPv6 address.

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# destination-ip-address 50.1.1.1  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```

duration

Use this command to define the duration of the EDPL test in hours, minutes, and seconds.

Use the *no* form of this command to reset the duration to its default.



Note: If duration is not configured, the EDPL test will run forever until it is stopped by using CLI command.

Command Syntax

```
duration hours <1-60> minutes <1-60> seconds <1-60>
no duration
```

Parameters

hours

Duration in hours (1–60).

minutes

Duration in minutes (1–60).

seconds

Duration in seconds (1–60).

Default

None.

Command Mode

EDPL Test Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl test test-1
OcNOS(edpl-test)# duration hours 1 minutes 10
OcNOS(edpl-test)# commit
OcNOS(edpl-test)# end
OcNOS#
```

edpl profile

Use this command to create or enter an Ethernet Data Plane Loopback (EDPL) profile configuration mode, where loopback parameters such as direction, address family, and filters can be defined.

Use the *no* form of this command to delete an existing EDPL profile.

Command Syntax

```
edpl profile <name>  
no edpl profile <name>
```

Parameters

<name>

Specifies the name of the EDPL profile.

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```


edpl start

Use this command to start an EDPL test based on the defined test configuration.

Once the test starts, packets matching the configured profile are looped back for the specified duration or until manually stopped.

Command Syntax

```
edpl start <name>
```

Parameters

<name>

Name of the EDPL test to start.

Default

None.

Command Mode

EDPL Test Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl start test-1
```

edpl stop

Use this command to stop a running EDPL test before the configured duration expires.

Command Syntax

```
edpl stop <name>
```

Parameters

<name>

Name of the EDPL test to stop.

Default

None.

Command Mode

EDPL Test Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl stop test-1
```

edpl test

Use this command to create or enter an EDPL (Ethernet Data Plane Loopback) test configuration mode, where you can associate an EDPL profile, define test duration, and specify the source interface for loopback.

Use the *no* form of this command to delete an existing EDPL test configuration.

Command Syntax

```
edpl test <name>  
no edpl test <name>
```

Parameters

<name>

Specifies the name of the EDPL test instance.

Default

None.

Command Mode

EDPL Test Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl test test-1  
OcNOS(edpl-test)# commit  
OcNOS(edpl-test)# end  
OcNOS#
```

ethertype

Use this command to configure the EtherType value for filtering Layer 2 frames in EDPL tests.

Use the *no* form of this command to remove the configured EtherType.

Command Syntax

```
ethertype <0xvalue>  
no ethertype
```

Parameters

<0xvalue>

Ethernet IP in hexadecimal format (for example, 0x800 for IPv4).

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# ethertype 0x800  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```

ip-swap

Use this command to configure EDPL to swap only IP addresses (source and destination) during packet loopback. Use the *no* form of this command to disable IP address swapping.

Command Syntax

```
ip-swap
no ip-swap
```

Parameters

None.

Default

Disabled.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1
OcNOS(edpl-profile-config)# ip-swap
OcNOS(edpl-profile-config)# commit
OcNOS(edpl-profile-config)# end
OcNOS#
```

source-port

Use this command to specify the source interface on which the EDPL test runs. If the source-port is a logical interface (SVI, IRB, or BVI), specify the L2 parent interface using the l2-port option.

Use the *no* form of this command to remove the configured source port.

Command Syntax

```
source-port <interface-name> [l2-port <interface-name>]  
no source-port <interface-name>
```

Parameters

<interface-name>

Name of the source physical, sub-interface, or logical interface.

l2-port <interface-name>

Specifies the parent L2 port for logical interfaces (SVI, IRB, BVI).

Default

None.

Command Mode

EDPL Test Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

For a physical interface:

```
OcNOS(config)# edpl test test-1  
OcNOS(edpl-test)# source-port xe-0/0/1  
OcNOS(edpl-test)# commit  
OcNOS(edpl-test)# end  
OcNOS#
```

For a logical interface (SVI):

```
OcNOS(config)# edpl test test-1  
OcNOS(edpl-test)# source-port vlan10 l2-port xe-0/0/2  
OcNOS(edpl-test)# commit  
OcNOS(edpl-test)# end  
OcNOS#
```

source-ip-address

Use this command to specify the source IP address for filtering IPv4 or IPv6 packets during EDPL tests.

Use the *no* form of this command to remove the configured source IP filter.

Command Syntax

```
source-ip-address <ip-address>  
no source-ip-address
```

Parameters

<ip-address>

Source IPv4 or IPv6 address.

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# source-ip-address 10.1.1.1  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```

source-udp-port

Use this command to configure the source UDP port number used for matching packets in EDPL tests.

Use the *no* form of this command to remove the configured source UDP port.

.

Command Syntax

```
source-udp-port <port-number>
no source-udp-port
```

Parameters

<port-number>

Source UDP port number (1–65535).

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1
OcNOS(edpl-profile-config)# source-udp-port 1000
OcNOS(edpl-profile-config)# commit
OcNOS(edpl-profile-config)# end
OcNOS#
```


udp-tcp-port-swap

Use this command to configure EDPL to swap IP addresses and UDP/TCP port numbers during packet loopback. Use the *no* form of this command to disable IP and port swapping.

Command Syntax

```
udp-tcp-port-swap
no udp-tcp-port-swap
```

Parameters

None.

Default

Disabled.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1
OcNOS(edpl-profile-config)# udp-tcp-port-swap
OcNOS(edpl-profile-config)# commit
OcNOS(edpl-profile-config)# end
OcNOS#
```

vlan

Use this command to configure the VLAN ID for matching Layer 2 packets during EDPL testing.

Use the *no* form of this command to remove the VLAN filter.

Command Syntax

```
vlan <vlan-id>  
no vlan
```

Parameters

<vlan-id>

VLAN ID value (1–4094).

Default

None.

Command Mode

EDPL Profile Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS(config)# edpl profile edpl-1  
OcNOS(edpl-profile-config)# vlan 10  
OcNOS(edpl-profile-config)# commit  
OcNOS(edpl-profile-config)# end  
OcNOS#
```

show ethernet edpl

Use this command to display all EDPL tests currently in progress on the device.

The output includes test name, direction, and remaining duration.

Command Syntax

```
show ethernet edpl
```

Parameters

None.

Default

None.

Command Mode

Execution mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS# show ethernet edpl
Name                Direction    Duration
-----
test-1              Upstream    01:00:00 (00:59:57 remaining)
OcNOS#
```

show ethernet edpl detail

Use this command to display detailed information for all EDPL tests, including applied filters, interfaces, actions (swap parameters), and test timers.

Command Syntax

```
show ethernet edpl detail
```

Parameters

None.

Default

None.

Command Mode

Execution mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Examples

```
OcNOS# show ethernet edpl detail
test-1 (xe-0/0/1)
-----
Direction      Upstream
Filters         ipv4-family
Action          ip-swap, udp-tcp-port-swap
Duration        01:00:00 (00:59:56 remaining)
OcNOS#
```

Troubleshooting

1. When EDPL traffic is not looping back as expected:

- Verify that the EDPL profile and test configuration are properly committed using the show running-config | include edpl command.
- Check if the correct direction (upstream or downstream) and address-family (L2, IPv4, IPv6) are configured in the profile.
- Confirm that udp-tcp-port-swap or ip-swap actions are enabled as per the test type.

2. If IPv4 or IPv6 packets are not reflected correctly:

- Check the source and destination IP address and UDP/TCP port filter configurations under the EDPL profile to ensure they match the test generator's traffic pattern.
- Verify that the source-port in the EDPL test corresponds to an active physical or logical interface.
- For logical interfaces (SVI/IRB/BVI), ensure the L2 parent port is correctly mapped using the I2-port option.

- Review the ingress and egress interface counters using show interface <port> counters to confirm packet reception and transmission.
3. When L2 loopback is not occurring properly:
- Check that the destination MAC, VLAN, and ethertype configured in the EDPL profile match the test traffic parameters.
 - Verify that the direction is correctly set to upstream or downstream based on the expected loopback path.
 - Use a packet capture tool on the host to confirm whether MAC swapping is performed correctly.
 - Ensure the interface used for the test is in an up/up state and is not part of another service that may filter L2 frames.
4. When EDPL test fails to start or stops prematurely:
- Confirm the test is properly associated with an existing profile using show ethernet edpl detail.
 - Check that the duration is within supported limits (1–60 hours/minutes/seconds).
 - Verify that no other EDPL tests are running concurrently on the same interface, as overlapping tests may conflict.
 - Review system logs with show log | include edpl for any test initialization or configuration errors.
 - Validate that the hardware-profile filter groups required for the configured address family (for example, egress-l2, egress-ipv4, or egress-ipv6) are enabled.
5. When host does not receive reflected packets:
- Ensure that unicast traffic is being sent, as EDPL does not support multicast or broadcast packets.
 - Check if the test direction is supported (for example, L3 reflection is only supported in the upstream direction).
 - Verify that the EDPL reflector device is reachable and properly configured to respond to host.
 - Review packet headers on the host to confirm IP/MAC/UDP swaps occur as configured in the profile.
6. To confirm EDPL operational status and counters:
- Use show ethernet edpl to confirm the test is running and monitor remaining duration.
 - Use show ethernet edpl detail to verify the active filters, actions, and associated source interface.
 - Check show interface <port> counters to verify that both RX and TX counters increment during the test, indicating traffic loopback activity.
 - If counters remain static, ensure that the interface and VLAN configurations align with the generator's traffic setup.
7. When performance metrics deviate or show inconsistencies:
- Verify that no intermediate devices are filtering or shaping the test traffic.
 - Check for CPU or hardware resource utilization on the host that might affect loopback performance.
 - If using IPv6, ensure only the first 64 bits of the address are used for ACL matching, as per platform limitations.
 - Repeat the test with reduced traffic rates to confirm hardware stability before full-rate testing.

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms or Acronym | Description |
|---------------------------------------|---|
| Ethernet Data Plane Loopback (EDPL) | A feature that enables loopback of Ethernet data-plane traffic on a device, allowing remote testing of interface performance based on RFC 2544 methodology. |
| VLAN | A Virtual Local Area Network identifier used to segregate Layer 2 broadcast domains. It can be used as a filtering parameter in L2 EDPL profiles. |
| MAC Address | The hardware address of a network interface used to filter L2 traffic during EDPL tests. |
| Switched Virtual Interface (SVI) | A logical Layer 3 interface associated with a VLAN, often used as a source-port in EDPL testing. |
| Integrated Routing and Bridging (IRB) | A logical interface that provides connectivity between Layer 2 and Layer 3 domains. Can participate in EDPL testing using the I2-port mapping. |
| Bridge Virtual Interface (BVI) | A logical interface representing a bridge group in routed bridge configurations, supported as an EDPL source-port. |
| Egress / Ingress | Egress refers to the outgoing interface where packets exit the device; ingress refers to the incoming interface where packets are received. |

LAYER 3 VIRTUAL PRIVATE NETWORK CONFIGURATION

| | |
|---|------|
| MPLS Layer-3 VPN Configurations | 2341 |
| Overview | 2341 |
| Configure MPLS Layer-3 VPN | 2343 |
| Verify the MPLS-VPN Configuration | 2355 |
| Inter-AS VPN Configuration Overview | 2359 |
| Types of Inter-AS VPN | 2359 |
| Inter-AS VPN Option-A Overview | 2359 |
| Inter-AS VPN Option-B Overview | 2370 |
| Inter-AS VPN Option-C Configuration | 2382 |
| ECMP Support for L3VPN | 2418 |
| Example 1 | 2432 |
| Example 2 | 2432 |
| ECMP Support for L3EVPN | 2433 |
| L3VPN over EVPN MPLS Configuration | 2498 |
| Topology | 2498 |
| Configuration | 2498 |
| L3VPN GR Configuration | 2522 |
| Topology | 2522 |
| Configuration | 2523 |
| Validation | 2527 |
| TWAMP over L3VPN Configuration | 2531 |
| Topology | 2531 |
| Configure TWAMP over L3VPN for LDP | 2532 |
| Validation | 2537 |
| Bridge Virtual Interface (BVI) Over L3VPN | 2540 |
| Overview | 2540 |
| Prerequisites | 2541 |
| Configuration | 2542 |
| Running Configuration | 2545 |
| Validation | 2553 |
| Glossary | 2554 |
| DSCP Preserve | 2555 |
| DSCP Preserve Characteristics | 2555 |
| DSCP Preserve Global | 2556 |
| Configuration for DSCP Preserve Global | 2556 |

| | |
|---|------|
| Commands | 2568 |
| DSCP Preserve Per-VRF | 2570 |
| DSCP Preserve Per-VRF Characteristics | 2570 |
| Configuration for DSCP Preserve Per-VRF | 2571 |
| DSCP Commands | 2586 |
| Glossary | 2588 |

Refer to the *DHCP Relay Agent Over L3VPN Configuration* section in *System Management Configuration guide*.

MPLS Layer-3 VPN Configurations

This chapter contains configurations for MPLS Layer-3 Virtual Private Networks (VPNs).

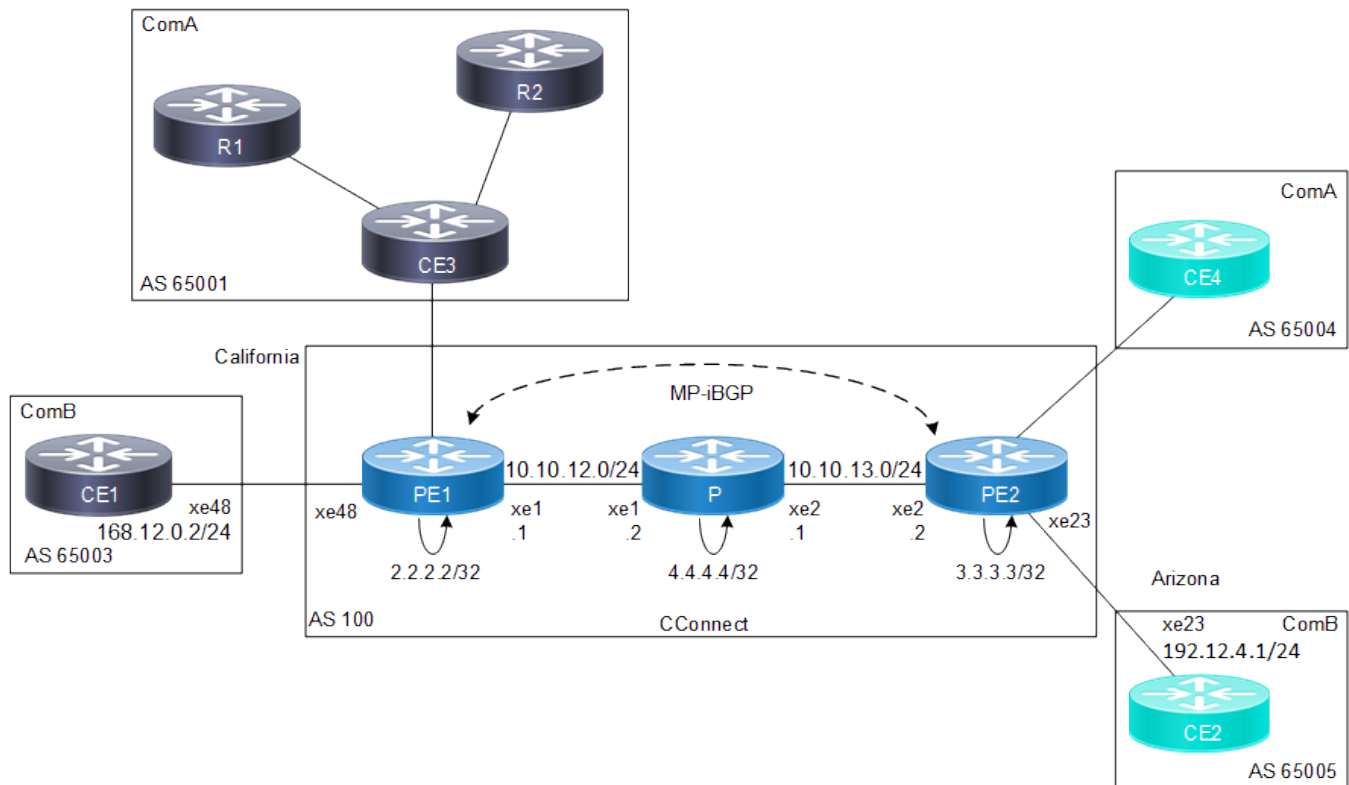
Overview

The MPLS Layer-3 VPN solution provides address space and routing separation via the use of per-VPN Routing and Forwarding tables (VRFs), and MPLS switching in the core and at the edge of the network. VPN customer routing data is imported into the VRFs utilizing the Route Target BGP extended community. This routing data is identified by a Route Distinguisher (RD) and is distributed among Provider Edge (PE) routers using Multi-Protocol BGP extensions.

Terminology

The following illustrates a Virtual Private Network in a CConnect Service Provider Network. This illustration corresponds to the terms defined in this subsection.

Figure 87. CConnect Provider with ComA and ComB Customers



- **Service Provider**

The organization that owns the infrastructure that provides leased lines to customers, offering them a Virtual Private Network Service. In the above illustration, CConnect is the service provider providing services to clients ComA and ComB.

- **Customer Edge (CE) Router**

A router at a customer's site that connects to the Service Provider via one or more Provider Edge routers. In the above illustration, CE1, CE2, CE3 and CE4 are all CE routers connected directly to the CConnect network.

- **Provider Edge (PE) Router**

A provider's router connected to a CE router through a leased line or dial-up connection. In the above illustration, PE1 and PE2 are the PE routers, because they link the CConnect service provider to its clients.

- **Provider Core Router (P)**

The devices in the core of the service provider network, which are generally not Provider Edge routers. In the above illustration, the P router is the Provider device, not connected to any customer, and is the core of the CConnect network.

- **Site**

A contiguous part of the customer network. A site connects to the provider network through transmission lines, using a CE and PE router. In the above illustration, R1, R2 and CE3 comprise a Customer network, and are seen as a single site by the CConnect network.

- **Customer Router**

In the illustration above, R1 and R2 are the Customer routers, and are not directly connected to the CConnect network.

VPN Routing Process

The OcNOS MPLS-VPN Routing process follows these steps:

1. Service Providers provide VPN services from PE routers that communicate directly with CE routers via an Ethernet Link.
2. Each PE router maintains a Routing and Forwarding table (VRF) for each customer. This guarantees isolation, and allows the usage of uncoordinated private addresses. When a packet is received from the CE, the VRF that is mapped to that site is used to determine the routing for the data. If a PE has multiple connections to the same site, a single VRF is mapped to all of those connections.
3. After the PE router learns of the IP prefix, it converts it into a VPN-IPv4 prefix by prepending it with an 8-byte Route Distinguisher (RD). The RD ensures that even if two customers have the same address, two separate routes to that address can be maintained. These VPN-IPv4 addresses are exchanged between the PE routers through MP-BGP.
4. A unique Router ID (usually the loopback address) is used to allocate a label, and enable VPN packet forwarding across the backbone.
5. Based on routing information stored in the VRF table, packets are forwarded to their destination using MPLS. Each PE router allocates a unique label to every route in each VRF (even if they have the same next hop), and propagates these labels, together with 12-byte VPN-IPv4 addresses, through Multi-Protocol BGP.
6. Ingress PE routers prepend a two-level label stack to the VPN packet, which is forwarded across the Provider network. This label stack contains a BGP-specific label from the VRF table (associated with the incoming interface), specifying the BGP next hop and an LDP-specific label from the global FTN table, specifying the IP next hop.
7. The Provider router in the network switches the VPN packet, based on the top label or the LDP-specific label in the stack. This top label is used as the key to lookup in the incoming interface's Incoming Labels Mapping table (ILM). If there is an outbound label, the label is swapped, and the packet is forwarded to the next hop; if not, the router is the penultimate router, and it pops the LDP-specific label, and forwards the packet with only the BGP-specific label to the egress PE router.

8. The egress PE router pops the BGP-specific label, performs a single label lookup in the outbound interface, and sends the packet to the appropriate CE router.

Configure MPLS Layer-3 VPN

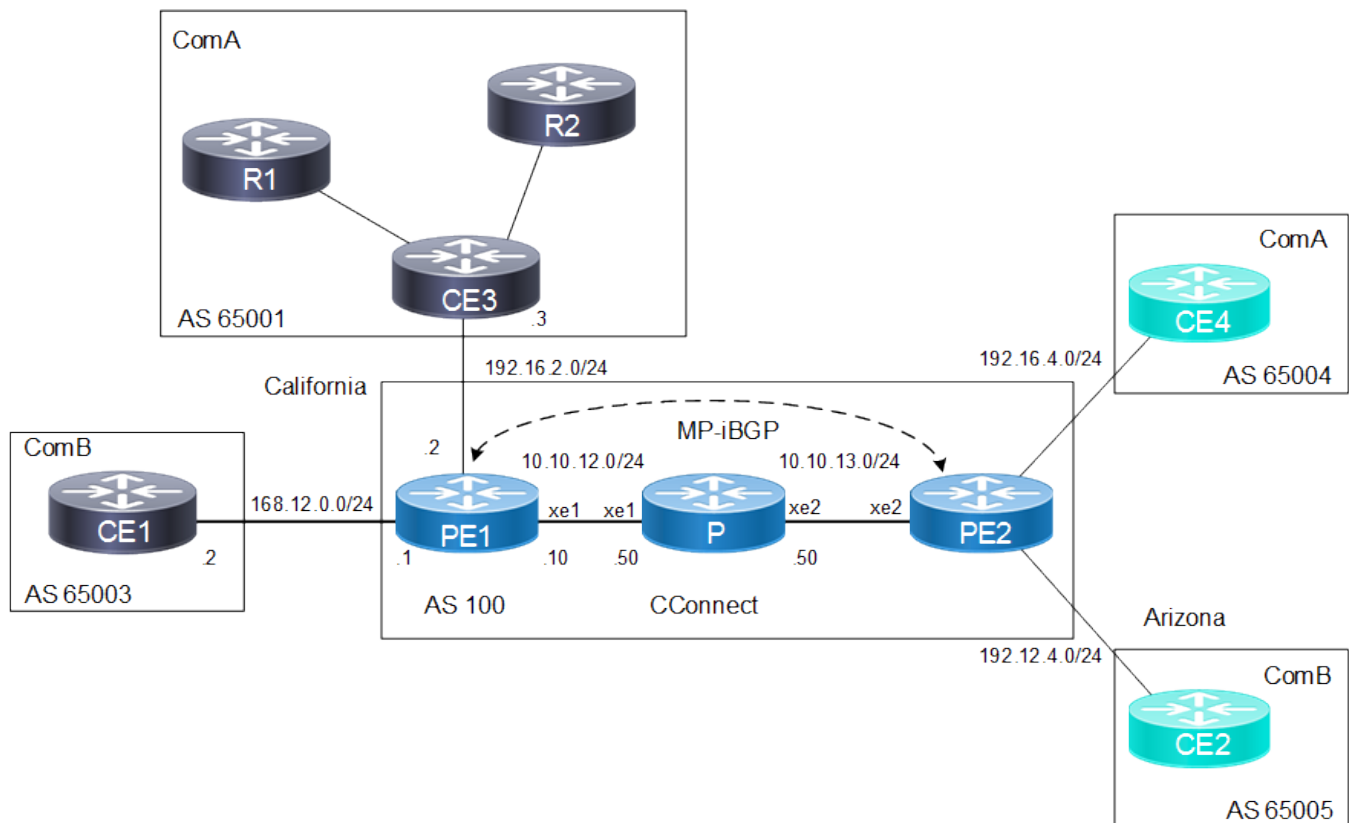
The MPLS Layer-3 VPN configuration process can be divided into the following tasks

1. Establish connection between PE routers.
2. Configure PE1 and PE2 as iBGP neighbors.
3. Create VRF.
4. Associate interfaces to VRFs.
5. Configure VRF Route Destination and Route Targets.
6. Configure CE neighbor for the VPN.
7. Verify the MPLS to VPN configuration.

Topology

In this example, the CConnect MPLS-VPN backbone has two customers — ComA and ComB. Both customers have sites in California and Arizona. The following topology shows BGP4 address assignment between PE and CE routers. The steps that follow provision a customer VPN service across the MPLS-VPN backbone.

Figure 88. Connect Sample Topology



Establishing this connection involves three steps:

Enable Label Switching

This is a sample configuration to enable label switching for the Labeled Switched Path (LSP) between PE1 and PE2.

Enable Label Switching: PE1

| | |
|---|---|
| configure terminal | Enter configure mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 2.2.2.2/32 secondary | Assign the IPv4 address |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe1 | Enter Interface mode |
| (config-if)#ip address 10.10.12.1/24 | Assign IPv4 address |
| (config-if)#label-switching | Enabling label switching capability on router |
| (config-if)#commit | Commit the transaction. |

Enable Label Switching: P

| | |
|--|---|
| configure terminal | Enter configure mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 4.4.4.4/32 area 0 secondary | Assign the IPv4 address |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe1 | Enter Interface mode |
| (config-if)#ip address 10.10.12.2/24 | Assign IPv4 address |
| (config-if)#label-switching | Enabling label switching capability on router |
| (config-if)#commit | Commit the transaction. |
| (config)#interface xe2 | Enter Interface mode |
| (config-if)#ip address 10.10.13.1/24 | Assign IPv4 address |
| (config-if)#label-switching | Enabling label switching capability on router |
| (config-if)#commit | Commit the transaction. |

Enable Label Switching: PE2

| | |
|---|---|
| configure terminal | Enter configure mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 3.3.3.3/32 secondary | Assign the IPv4 address |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe2 | Enter Interface mode |
| (config-if)#ip address 10.10.13.2/24 | Assign IPv4 address |
| (config-if)#label-switching | Enabling label switching capability on router |
| (config-if)#commit | Commit the transaction. |

Enable IGP

This is a sample configuration to establish connections between the two Provider Edge routers PE1 and PE2.

Enable IGP Switching: PE1

| | |
|---|--|
| <code>configure terminal</code> | Enter configure mode |
| <code>(config)#router ospf 100</code> | Configure the routing process and specify the Process ID (100) |
| <code>(config-router)#network 10.10.12.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#network 2.2.2.2/32 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#commit</code> | Commit the transaction. |

Enable IGP Switching: P

| | |
|---|--|
| <code>configure terminal</code> | Enter configure mode |
| <code>(config)#router ospf 100</code> | Configure the routing process and specify the Process ID (100) |
| <code>(config-router)#network 10.10.12.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#network 10.10.13.0/24 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#network 4.4.4.4/32 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#commit</code> | Commit the transaction. |

Enable IGP Switching: PE2

| | |
|--|--|
| <code>configure terminal</code> | Enter configure mode |
| <code>(config)#router ospf 100</code> | Configure the routing process and specify the Process ID (100) |
| <code>(config-router)#network 3.3.3.3/32 area 0</code> | Define the interface on which OSPF runs and associate the area ID (0) with the interface |
| <code>(config-router)#commit</code> | Commit the transaction. |

Enable Label Switching Protocol

Label switching protocols are used to set up a Label-Switched Path (LSP) between PE routers. OcNOS supports LDP and RSVP-TE protocols for label switching. Enable either LDP or RSVP-TE.

This is a sample configuration to enable LDP on the whole path between PE1 and PE2.

LDP: PE1

| | |
|-----------------------------|---------------------------|
| configure terminal | Enter configure mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#exit | Exit router mode |
| (config)#interface xe1 | Enter interface mode |
| (config-if)#enable-ldp ipv4 | Enabling LDP on interface |
| (config-if)#commit | Commit the transaction. |

LDP: P

| | |
|-----------------------------|---------------------------|
| configure terminal | Enter configure mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#exit | Exit router mode |
| (config)#interface xe1 | Enter interface mode |
| (config-if)#enable-ldp ipv4 | Enabling LDP on interface |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe2 | Enter interface mode |
| (config-if)#enable-ldp ipv4 | Enabling LDP on interface |
| (config-if)#commit | Commit the transaction. |

LDP: PE2

| | |
|-----------------------------|---------------------------|
| configure terminal | Enter configure mode |
| (config)#router ldp | Enter router mode for LDP |
| (config-router)#exit | Exit router mode |
| (config)#interface xe2 | Enter interface mode |
| (config-if)#enable-ldp ipv4 | Enabling LDP on interface |
| (config-if)#commit | Commit the transaction. |

This is a sample configuration to enable RSVP-TE along the entire path between PE1 and PE2.

RSVP-TE: PE1

| | |
|-----------------------------------|----------------------------------|
| configure terminal | Enter configure mode |
| (config)#router rsvp | Enter Configure Router mode |
| (config-router)#rsvp-path p1 mpls | Enter the path mode for RSVP P1. |
| (config-path)#10.10.12.2 loose | Configure loose path |
| (config-path)#exit | Exit Configure Router mode |
| (config)#rsvp-trunk t1 | Configure RSVP trunk t1 |

| | |
|-------------------------------|---|
| (config-rsvp)#primary path p1 | Specify an RSVP path to be used |
| (config-rsvp)#from 2.2.2.2 | Assign the source loopback address to the RSVP trunk |
| (config-rsvp)#to 3.3.3.3 | Assign the source loopback address to the to the RSVP trunk |
| (config-rsvp)#exit | Exit RSVP trunk mode |
| (config)#interface xe1 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP in interface |
| (config-if)#commit | Commit the transaction. |

RSVP-TE: P

| | |
|-------------------------|-----------------------------|
| configure terminal | Enter configure mode |
| (config)#router rsvp | Enter Configure Router mode |
| (config-router)#exit | Exit Configure Router mode |
| (config)#interface xe1 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP in interface |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe2 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP in interface |
| (config-if)#commit | Commit the transaction. |

RSVP-TE: PE2

| | |
|-------------------------------|---|
| configure terminal | Enter configure mode |
| (config)#router rsvp | Enter Configure Router mode |
| (config-router)#rsvp-trunk t1 | Configure RSVP trunk t1 |
| (config-rsvp)#from 3.3.3.3 | Assign the source loopback address to the RSVP trunk |
| (config-rsvp)#to 2.2.2.2 | Assign the source loopback address to the to the RSVP trunk |
| (config-rsvp)#exit | Exit RSVP trunk mode |
| (config)#interface xe2 | Enter the interface mode |
| (config-if)#enable-rsvp | Enable RSVP in interface |
| (config-if)#commit | Commit the transaction. |

Configure PEs as BGP Neighbors

BGP is the preferred protocol to transport VPN routes because of its multiprotocol capability and its scalability. Its ability to exchange information between indirectly connected routers supports keeping VPN routing information out of the Provider (P) routers. The P routers carry information as an optional BGP attribute. Additional attributes are transparently forwarded by any P router. The MPLS-VPN forwarding model does not require the P routers to make

routing decisions based on VPN addresses: They forward packets based on the label value attached to the packet. The P routers do not require a VPN configuration in order to carry this information.

PE1

| | |
|--|--|
| #configure terminal | Enter the configure terminal |
| (config)#router bgp 100 | Enter the Router BGP mode, ASN: 100 |
| (config-router)#bgp router-id 2.2.2.2 | Router identifier for BGP |
| (config-router)#exit | Exit router BGP mode |
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#exit | Exit the Router LDP mode |
| (config)#router bgp 100 | Enter the Router BGP mode, ASN: 100 |
| (config-router)#neighbor 3.3.3.3 remote-as 100 | Configuring ABR1 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 3.3.3.3 update-source 2.2.2.2 | Source of routing updates |
| (config-router)#address-family vpnv4 unicast | Configure VPNv4 address family |
| (config-router-af)#neighbor 3.3.3.3 activate | Activate the VPN neighbor |
| (config-router-af)#exit-address-family | Exit address family mode |
| (config-router)#address-family ipv4 unicast | Configure IPv4 address family |
| (config-router-af)#neighbor 3.3.3.3 activate | Activate the IPv4 neighbor |
| (config-router-af)#commit | Commit the transaction. |

PE2

| | |
|--|--|
| #configure terminal | Enter the configure terminal |
| (config)#router bgp 100 | Enter the Router BGP mode, ASN: 100 |
| (config-router)#bgp router-id 3.3.3.3 | Router identifier for BGP |
| (config-router)#exit | Exit router BGP mode |
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#exit | Exit the Router LDP mode |
| (config)#router bgp 100 | Enter the Router BGP mode, ASN: 100 |
| (config-router)#neighbor 2.2.2.2 remote-as 100 | Configuring ABR1 as iBGP neighbor using it's loopback IP |
| (config-router)#neighbor 2.2.2.2 update-source 3.3.3.3 | Source of routing updates |
| (config-router)#address-family vpnv4 unicast | Configure VPNv4 address family |
| (config-router-af)#neighbor 2.2.2.2 activate | Activate the VPN neighbor |
| (config-router-af)#exit-address-family | Exit address family mode |
| (config-router)#address-family ipv4 unicast | Configure IPv4 address family |

| | |
|---|----------------------------|
| (config-router-af) #neighbor 2.2.2.2 activate | Activate the IPv4 neighbor |
| (config-router-af) #commit | Commit the transaction. |

Create VRF

Each PE router in the MPLS-VPN backbone is attached to a site that receives routes from a specific VPN, so the PE router must have the relevant Virtual Routing and Forwarding (VRF) configuration for that VPN.

This command creates a VRF RIB (Routing Information Base), assigns a VRF-ID, and switches the command mode to `vrf` mode. The following example creates a VRF named `ComB`.

PE1

| | |
|-----------------------|--------------------------------------|
| #configure terminal | Enter the configure terminal |
| (config) #ip vrf ComB | Configure VRF instance |
| (config-vrf) #rd 2:1 | Configure Router Distinguisher value |
| (config-vrf) #exit | Exit VRF mode |
| (config) #ip vrf ComA | Configure VRF instance |
| (config-vrf) #rd 1:2 | Configure Router Distinguisher value |
| (config-vrf) #exit | Exit VRF mode |
| (config) #commit | Commit the transaction. |

PE2

| | |
|-----------------------|--------------------------------------|
| #configure terminal | Enter the configure terminal |
| (config) #ip vrf ComB | Configure VRF instance |
| (config-vrf) #rd 1:2 | Configure Router Distinguisher value |
| (config-vrf) #exit | Exit VRF mode |
| (config) #ip vrf ComA | Configure VRF instance |
| (config-vrf) #rd 2:1 | Configure Router Distinguisher value |
| (config-vrf) #exit | Exit VRF mode |
| (config) #commit | Commit the transaction. |

Associate Interfaces to VRFs

After the VRFs are defined on the PE router, the PE router needs to recognize which interfaces belong to which VRF. The VRF is populated with routes from connected sites. More than one interface can belong to the same VRF.

In the following example, interface `xe48` is associated with the VRF named `ComB`.

PE1

| | |
|--------------------------|------------------------------|
| #configure terminal | Enter the configure terminal |
| (config) #interface xe49 | Enter interface mode |

| | |
|---------------------------------------|--|
| (config-if)#ip vrf forwarding ComA | Bind the VRF instance to the interface |
| (config-if)exit | Exit interface mode |
| (config)#interface xe48 | Enter interface mode |
| (config-if)#ip vrf forwarding ComB | Bind the VRF instance to the interface |
| (config-if)#ip address 168.12.0.3/24 | Assign IPv4 address |
| PE1(config-if)#ip vrf forwarding ComA | Bind the VRF instance to the interface |
| (config-if)exit | Exit interface mode |
| (config)#commit | Commit the transaction. |

PE2

| | |
|--------------------------------------|--|
| #configure terminal | Enter the configure terminal |
| (config)#interface xe48 | Enter interface mode |
| (config-if)#ip vrf forwarding ComA | Bind the VRF instance to the interface |
| (config-if)#ip address 192.16.4.2/24 | Assign IPv4 address |
| (config-if)exit | Exit interface mode |
| (config)#interface xe49 | Enter interface mode |
| (config-if)#ip vrf forwarding ComB | Bind the VRF instance to the interface |
| (config-if)#ip address 168.12.4.2/24 | Assign IPv4 address |
| (config-if)exit | Exit interface mode |
| (config)#commit | Commit the transaction. |

Configure VRF—RD and Route Targets

After the VRF is created, configure Router Distinguishers and the Route Targets.

Configure Route Distinguishers

Route Distinguishers (RDs) make all customer routes unique. The routes must be unique, so that Multi-Protocol BGP treats the same prefix from two different VPNs as non-comparable routes. To configure RDs, a sequence of 64 bits is prepended to the IPv4 address in the Multi-Protocol BGP update. BGP considers two IPv4 addresses with different RDs as non-comparable, even if they have the same address and mask.

Assign a particular value to the RD for each VRF on the PE router. To display the routing table for a VRF, use the `show ip route vrf` command.

The following example shows adding an RD. Configure a VRF in both PEs with a unique RD value:

PE1

| | |
|--------------------------------------|-------------------------------------|
| #configure terminal | Enter the configure terminal |
| (config)#ip vrf ComA | Configure VRF instance |
| (config-vrf)#rd 2:2 | Configure Router Distiguisher value |
| (config-vrf)#route-target both 200:1 | Configure route-target as both |

| | |
|--------------------------------------|--------------------------------------|
| (config-vrf)#exit | Exit VRF mode |
| (config)#ip vrf ComB | Configure VRF instance |
| (config-vrf)#rd 1:1 | Configure Router Distinguisher value |
| (config-vrf)#route-target both 100:1 | Configure route-target as both |
| (config-vrf)#exit | Exit VRF mode |
| (config)#commit | Commit the transaction. |

PE2

| | |
|--------------------------------------|--------------------------------------|
| #configure terminal | Enter the configure terminal |
| (config)#ip vrf ComA | Configure VRF instance |
| (config-vrf)#rd 2:1 | Configure Router Distinguisher value |
| (config-vrf)#route-target both 200:1 | Configure route-target as both |
| (config-vrf)#exit | Exit VRF mode |
| (config)#ip vrf ComB | Configure VRF instance |
| (config-vrf)#rd 1:2 | Configure Router Distinguisher value |
| (config-vrf)#route-target both 100:1 | Configure route-target as both |
| (config-vrf)#exit | Exit VRF mode |
| (config)#commit | Commit the transaction. |

Configure Route Targets

Any routes learned from customers are advertised across the network through Multi-Protocol BGP, and any routes learned through Multi-Protocol BGP are added into the appropriate VRFs. The route target helps PE routers identify which VRFs should receive the routes.

The `route-target` command creates lists of import and export route-target extended communities for the VRF. It specifies a target VPN extended community. Execute the command once for each community. All routes with the specific route-target extended community are imported into all VRFs with the same extended community as an import route-target.

The following example demonstrates the route-target configuration.

PE1

| | |
|--------------------------------------|--------------------------------------|
| #configure terminal | Enter the configure terminal |
| (config)#ip vrf ComA | Configure VRF instance |
| (config-vrf)#rd 2:2 | Configure Router Distinguisher value |
| (config-vrf)#route-target both 200:1 | Configure route-target as both |
| (config-vrf)#exit | Exit VRF mode |
| (config)#ip vrf ComB | Configure VRF instance |
| (config-vrf)#rd 1:1 | Configure Router Distinguisher value |
| (config-vrf)#route-target both 100:1 | Configure route-target as both |

| | |
|--------------------------------|-------------------------|
| <code>(config-vrf)#exit</code> | Exit VRF mode |
| <code>(config)#commit</code> | Commit the transaction. |

PE2

| | |
|---|--------------------------------------|
| <code>#configure terminal</code> | Enter the configure terminal |
| <code>(config)#ip vrf ComA</code> | Configure VRF instance |
| <code>(config-vrf)#rd 2:1</code> | Configure Router Distinguisher value |
| <code>(config-vrf)#route-target both 200:1</code> | Configure route-target as both |
| <code>(config-vrf)#exit</code> | Exit VRF mode |
| <code>(config)#ip vrf ComB</code> | Configure VRF instance |
| <code>(config-vrf)#rd 1:2</code> | Configure Router Distinguisher value |
| <code>(config-vrf)#route-target both 100:1</code> | Configure route-target as both |
| <code>(config-vrf)#exit</code> | Exit VRF mode |
| <code>(config)#commit</code> | Commit the transaction. |

Configure CE Neighbor for the VPN (Using BGP/ OSPF)

To provide a VPN service, the PE-router must be configured so that any routing information learned from a VPN customer interface can be associated with a particular VRF. This is achieved using any standard routing protocol process (OSPF, BGP or static routes etc). Use any one of the following configurations (BGP, or OSPF) to configure the CE neighbor.

Using BGP

The BGP sessions between PE and CE routers can carry different types of routes (VPN-IPv4, IPv4 routes). Address families are used to control the type of BGP session. Configure a BGP address family for each VRF on the PE-router, and a separate address family to carry VPN-IPv4 routes between PE routers. All non-VPN BGP neighbors are defined using the `IPv4 address mode`. Each VPN BGP neighbor is defined under its associated address family mode.

A separate address family entry is used for every VRF, and each address family entry can have multiple CE routers within the VRF.

The PE and CE routers must be directly connected for BGP4 sessions; BGP multihop is not supported between PE and CE routers.

The following example places the router in address family mode, and specifies company names, `ComA` and `ComB`, as the names of the VRF instance to associate with subsequent IPv4 address family configuration mode commands. This configuration is used when BGP is used for PE and CE.

PE1

| | |
|--|---------------------------|
| <code>configure terminal</code> | Enter configure mode |
| <code>(config)#router bgp 100</code> | Enter BGP router mode |
| <code>(config-router)#address-family ipv4 unicast</code> | Enter address family mode |

| | |
|--|--|
| <code>(config-router-af)#redistribute connected</code> | Redistribute connected addresses |
| <code>(config-router-af)#exit-address-family</code> | Exit address family mode |
| <code>(config-router)#address-family ipv4 vrf ComA</code> | Enter the IPv4 address family for VRF comA |
| <code>(config-router)#neighbor 192.16.2.3 remote-as 65001</code> | Specify the neighbor |
| <code>(config-router-af)#neighbor 192.16.2.3 activate</code> | Activate the neighbor |
| <code>(config-router-af)#exit-address-family</code> | Exit address family mode |
| <code>(config-router)#address-family ipv4 vrf ComB</code> | Enter the IPv4 address family for VRF comA |
| <code>(config-router)#neighbor 168.12.0.2 remote-as 65003</code> | Specify the neighbor |
| <code>(config-router-af)#neighbor 192.12.0.2 activate</code> | Activate the neighbor |
| <code>(config-router-af)#exit-address-family</code> | Exit address family mode |
| <code>(config-router)#commit</code> | Commit the transaction |

PE2

| | |
|--|--|
| <code>configure terminal</code> | Enter configure mode |
| <code>(config)#router bgp 100</code> | Enter BGP router mode |
| <code>(config-router)#address-family ipv4 unicast</code> | Enter address family mode |
| <code>(config-router-af)#redistribute connected</code> | Redistribute connected addresses |
| <code>(config-router-af)#exit-address-family</code> | Exit address family mode |
| <code>(config-router)#address-family ipv4 vrf ComA</code> | Enter the IPv4 address family for VRF comA |
| <code>(config-router)#neighbor 192.16.4.3 remote-as 65004</code> | Specify the neighbor |
| <code>(config-router-af)#neighbor 192.16.4.3 activate</code> | Activate the neighbor |
| <code>(config-router-af)#exit-address-family</code> | Exit address family mode |
| <code>(config-router)#address-family ipv4 vrf ComB</code> | Enter the IPv4 address family for VRF comA |
| <code>(config-router)#neighbor 168.12.4.1 remote-as 65005</code> | Specify the neighbor |
| <code>(config-router-af)#neighbor 192.12.4.1 activate</code> | Activate the neighbor |
| <code>(config-router-af)#exit-address-family</code> | Exit address family mode |
| <code>(config-router)#commit</code> | Commit the transaction |

CE1: BGP

| | |
|--|-------------------------------|
| <code>configure terminal</code> | Enter configure mode |
| <code>(config)#interface xe48</code> | Enter interface mode for xe48 |
| <code>(config-if)#ip address 168.12.0.2/24</code> | Assign IP address |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#router bgp 65003</code> | Enter BGP router mode |
| <code>(config-router)#neighbor 168.12.0.3 remote-as 100</code> | Specify the neighbor |
| <code>(config-router)#address-family ipv4 unicast</code> | Enter address family mode |

| | |
|--|----------------------------------|
| <code>(config-router-af)#neighbor 168.12.0.3 activate</code> | Activate the neighbor |
| <code>(config-router-af)#redistribute connected</code> | Redistribute connected addresses |
| <code>(config-router-af)#commit</code> | Commit the transaction |

CE2: BGP

| | |
|--|----------------------------------|
| <code>configure terminal</code> | Enter configure mode |
| <code>(config)#interface xe23</code> | Enter interface mode for xe48 |
| <code>(config-if)#ip address 192.12.4.1/24</code> | Assign IP address |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#router bgp 65005</code> | Enter BGP router mode |
| <code>(config-router)#neighbor 192.12.4.2 remote-as 100</code> | Specify the neighbor |
| <code>(config-router)#address-family ipv4 unicast</code> | Enter address family mode |
| <code>(config-router-af)#neighbor 192.12.4.2 activate</code> | Activate the neighbor |
| <code>(config-router-af)#redistribute connected</code> | Redistribute connected addresses |
| <code>(config-router-af)#commit</code> | Commit the transaction |

Using OSPF

Unlike BGP, OSPF does not run different routing contexts within one process. Thus, for running OSPF between the PE and CE routers, configure a separate OSPF process for each VRF that receives VPN routes through OSPF. The PE router distinguishes routers belonging to a specific VRF, by associating a particular customer interface to a specific VRF and to a particular OSPF process.

To redistribute VRF OSPF routes into BGP, redistribute OSPF under the BGP VRF address family submode.

PE1

| | |
|---|---|
| <code>configure terminal</code> | Enter configure mode |
| <code>(config)#router ospf 101 comA</code> | Enter OSPF router mode |
| <code>(config-router)#network 192.16.3.0/24 area 0</code> | Define the network on which OSPF runs and associate area ID |
| <code>(config-router)#redistribute bgp</code> | Redistribute BGP |
| <code>(config-router)#exit</code> | Exit router mode |
| <code>(config)#router ospf 102 comB</code> | Enter OSPF router mode |
| <code>(config-router)#network 168.12.0.2/24 area 0</code> | Define the network on which OSPF runs and associate area ID |
| <code>(config-router)#redistribute bgp</code> | Redistribute BGP |
| <code>(config-router)#commit</code> | Commit the transaction |
| <code>(config-router)#exit</code> | Exit router mode |
| <code>(config)#router bgp 100</code> | Enter the Router BGP mode, ASN: 100 |
| <code>(config-router)#address-family ipv4 vrf ComA</code> | Configure VRF address family |

| | |
|--|------------------------------|
| (config-router-af)#redistribute ospf | Redistribute OSPF |
| (config-router-af)#exit-address-family | Exit address family mode |
| (config-router)#address-family ipv4 vrf ComB | Configure VRF address family |
| (config-router-af)#redistribute ospf | Redistribute OSPF |
| (config-router-af)#redistribute rip | Redistribute RIP |
| (config-router-af)#commit | Commit the transaction |

PE2

| | |
|--|---|
| configure terminal | Enter configure mode |
| (config)#router ospf 101 comA | Enter OSPF router mode |
| (config-router)#network 192.16.4.0/24 area 0 | Define the network on which OSPF runs and associate area ID |
| (config-router)#redistribute bgp | Redistribute BGP |
| (config-router)#exit | Exit router mode |
| (config)#router ospf 102 comB | Enter OSPF router mode |
| (config-router)#network 168.12.0.3/24 area 0 | Define the network on which OSPF runs and associate area ID |
| (config-router)#redistribute bgp | Redistribute BGP |
| (config-router)#commit | Commit the transaction |
| (config-router)#exit | Exit router mode |
| (config)#router bgp 100 | Enter the Router BGP mode, ASN: 100 |
| (config-router)#address-family ipv4 vrf ComA | Configure VRF address family |
| (config-router-af)#redistribute ospf | Redistribute OSPF |
| (config-router-af)#exit-address-family | Exit address family mode |
| (config-router)#address-family ipv4 vrf ComB | Configure VRF address family |
| (config-router-af)#redistribute ospf | Redistribute OSPF |
| (config-router-af)#redistribute rip | Redistribute RIP |
| (config-router-af)#commit | Commit the transaction |

Verify the MPLS-VPN Configuration

Use the `show ip bgp neighbor` command to validate the neighbor session between the CE and the PE routers. Use the `show ip bgp vpnv4 all` command to display all the VRFs and the routes associated with them. The following is sample output for the above commands for the PE1, CE1 and PE2 routers (based on the topology illustrated above).

```
PE1#show ip bgp neighbors
BGP neighbor is 3.3.3.3, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 2.2.2.2, remote router ID 3.3.3.3
  BGP state = Established, up for 00:05:09
  Last read 00:00:13, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
```

```
Route refresh: advertised and received (old and new)
Address family IPv4 Unicast: advertised and received
Address family VPNv4 Unicast: advertised and received
Received 194 messages, 2 notifications, 0 in queue
Sent 198 messages, 3 notifications, 0 in queue
Route refresh request: received 2, sent 1
Minimum time between advertisement runs is 5 seconds
Update source is 2.2.2.2
For address family: IPv4 Unicast
BGP table version 8, neighbor version 8
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (both)
3 accepted prefixes
3 announced prefixes

For address family: VPNv4 Unicast
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (both)
0 accepted prefixes
2 announced prefixes

Connections established 5; dropped 4
Local host: 2.2.2.2, Local port: 35983
Foreign host: 3.3.3.3, Foreign port: 179
Nexthop: 2.2.2.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:05:14, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

BGP neighbor is 168.12.0.2, vrf ComB, remote AS 65003, local AS 100, external link
BGP version 4, local router ID 168.12.0.3, remote router ID 10.12.65.206
BGP state = Established, up for 00:34:38
Last read 00:00:14, hold time is 90, keepalive interval is 30 seconds
Neighbor capabilities:
Route refresh: advertised and received (old and new)
Address family IPv4 Unicast: advertised and received
Received 85 messages, 0 notifications, 0 in queue
Sent 86 messages, 0 notifications, 0 in queue
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 30 seconds
For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (standard)
2 accepted prefixes
0 announced prefixes

Connections established 1; dropped 0
Local host: 168.12.0.3, Local port: 179
Foreign host: 168.12.0.2, Foreign port: 36580
Nexthop: 168.12.0.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

CE1#show ip bgp neighbors
BGP neighbor is 168.12.0.3, remote AS 100, local AS 65003, external link
BGP version 4, local router ID 10.12.65.206, remote router ID 168.12.0.3
BGP state = Established, up for 00:36:14
Last read 00:00:10, hold time is 90, keepalive interval is 30 seconds
Neighbor capabilities:
Route refresh: advertised and received (old and new)
Address family IPv4 Unicast: advertised and received
Received 86 messages, 0 notifications, 0 in queue
```



```

Sent 89 messages, 0 notifications, 0 in queue
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 30 seconds
For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (both)
0 accepted prefixes
2 announced prefixes

Connections established 1; dropped 0
Local host: 168.12.0.2, Local port: 36580
Foreign host: 168.12.0.3, Foreign port: 179
Nexthop: 168.12.0.2
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

PE1#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, l -
labeled
                S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

  Network          Next Hop          Metric    LocPrf    Weight Path
Route Distinguisher: 1:1 (Default for VRF ComB)
*> 1 168.12.0.0/24    0.0.0.0              1         100      32768  ?
                        168.12.0.2          0         100      0 65003  ?
*>i 192.12.4.0        3.3.3.3              0         100      0 65005  ?
  Announced routes count = 2
  Accepted routes count = 1
Route Distinguisher: 1:2
*>i 192.12.4.0        3.3.3.3              0         100      0 65005  ?
  Announced routes count = 0
  Accepted routes count = 1
Route Distinguisher: 2:2 (Default for VRF ComA)
*> 1 192.16.2.0        0.0.0.0              1         100      32768  ?
                        192.16.2.3          0         100      0 65001  ?
  Announced routes count = 2
  Accepted routes count = 0

PE1#show ip bgp vpnv4 all neighbors 3.3.3.3 routes

For address family: VPNv4 Unicast
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, l -
labeled
                S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

  Network          Next Hop          Metric    LocPrf    Weight Path
Route Distinguisher: 1:2
*>i 192.12.4.0        3.3.3.3              0         100      0 65005  ?
  Accepted routes count = 1

PE2#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
                S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

  Network          Next Hop          Metric    LocPrf    Weight Path
Route Distinguisher: 1:2 (Default for VRF ComB)
*>i 168.12.0.0/24    2.2.2.2              0         100      0  ?
*> 1 192.12.4.0        0.0.0.0              0         100      32768  ?
  Announced routes count = 1
  Accepted routes count = 1
Route Distinguisher: 1:2
*>i 168.12.0.0/24    2.2.2.2              0         100      0  ?
  Announced routes count = 0
  Accepted routes count = 1

```

Verify MPLS-L3 VPN VRF Ping and Traceroute

Use the `ping mpls l3vpn` command for the below requirements:

- PE to PE L3VPN ping via VRF
- PE to remote CE Ping via the VRF
- CE to remote PE ping (to the VRF interface facing its customer edge).
- Trace route from PE to PE via VRF
- Trace route from PE to remote CE via VRF
- Commands for ipv6 ping and trace route

1. PE to PE L3VPN Ping via VRF:

```
PE2#ping 168.12.0.1 vrf ComB
Press CTRL+C to exit
PING 168.12.0.1 (168.12.0.1) 56(84) bytes of data.
64 bytes from 168.12.0.1: icmp_seq=1 ttl=64 time=0.695 ms

#
```

2. PE to remote CE Ping via VRF:

```
PE2#ping 168.12.0.2 vrf ComB
Press CTRL+C to exit
PING 168.12.0.2 (168.12.0.2) 56(84) bytes of data.
64 bytes from 168.12.0.2: icmp_seq=1 ttl=63 time=0.776 ms

--- 168.12.0.2 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.776/0.776/0.776/0.000 ms
PE2#
PE2
```

3. CE to remote PE ping:

```
CE1#ping 168.12.0.1
Press CTRL+C to exit
PING 168.12.0.1 (168.12.0.1) 56(84) bytes of data.
64 bytes from 168.12.0.1: icmp_seq=160 ttl=254 time=0.606 ms
64 bytes from 168.12.0.1: icmp_seq=161 ttl=254 time=0.558 ms
64 bytes from 168.12.0.1: icmp_seq=162 ttl=254 time=0.568 ms
64 bytes from 168.12.0.1: icmp_seq=163 ttl=254 time=0.574 ms
64 bytes from 168.12.0.1: icmp_seq=164 ttl=254 time=0.609 ms

--- 168.12.0.2 ping statistics ---
5 packets transmitted, 5 received, 0 errors, 0% packet loss, time 163002ms
```

4. Trace Route from PE to PE via VRF

```
PE2#traceroute ip 168.12.0.1 vrf ComB
traceroute to 168.12.0.1 (168.12.0.1), 30 hops max, 60 byte packets
1  168.12.0.1 (168.12.0.1)  0.706 ms  0.743 ms  0.989 ms
```

5. Trace Route from PE to Remote CE via VRF

```
PE2#traceroute ip 168.12.0.2 vrf ComB
traceroute to 168.12.0.2 (168.12.0.2), 30 hops max, 60 byte packets
1  168.12.0.1 (168.12.0.1)  0.871 ms  1.006 ms  1.055 ms
2  168.12.0.2 (168.12.0.2)  1.965 ms  2.045 ms  2.256 ms
```

Inter-AS VPN Configuration Overview

MPLS VPN architecture typically runs within an AS. Routes of any VPN can be flooded within the AS, but not to other ASs. To implement the exchange of VPN routes between different ASs, the inter-AS MPLS VPN model is used. The inter-AS MPLS VPN model is an extension to MPLS VPN framework. Route prefixes and labels can be advertised over links between different carrier networks through the inter-AS MPLS model.

The MPLS VPN solution serves an increasing number of users across many applications. A site at one geographical location often needs to connect to an ISP network at another geographical location. In this situation, for example, inter-AS issues may arise for operators who manage different metropolitan area networks (MANs) or backbone networks that span different autonomous systems (AS).

Types of Inter-AS VPN

1. Inter-AS VPN Option A: Autonomous system boundary routers (ASBRs) manage VPN routes for in-ter-AS VPNs through dedicated interfaces.
2. Inter-AS VPN Option B: ASBRs advertise labeled VPN-IPv4 routes to each other through MP-EBGP.
3. Inter-AS VPN Option C: PE devices advertise labeled VPN-IPv4 routes to each other through Mul-ti-hop MP-EBGP.

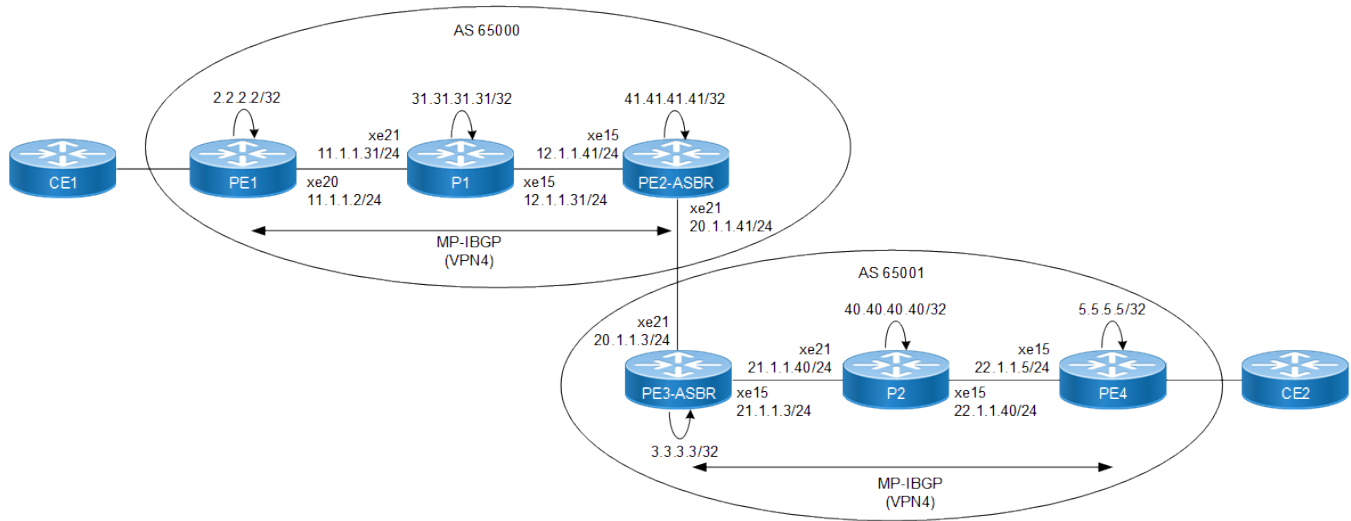
Inter-AS VPN Option-A Overview

This chapter explain about Inter-AS VPN Option-A. Option A is the simplest of the options to inter-connect the ASBRs Option A has the following characteristics:

- Each customer VRF requires either a physical interface or more likely a subinterface.
- Each ASBR thinks the other is a CE.
- One logical interface per VPN.
- Link may use any supported PE-CE protocol.
- Packets are sent unlabelled between the ASBRs.
- The most secure and easy option to provision.
- Does not scale well to a large number of VPNs.

Topology

Figure 89. InterAS-VPN Option-A



Inter-AS VPN Option-A Configuration

PE1

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 2.2.2.2/32 secondary | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#ip vrf vrf1 | Create a new VRF named vrf1. |
| (config-vrf)#rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| (config-vrf)#route-target both 1:1 | Import routes between route target (RT) ext-communities. |
| (config-vrf)#exit | Exit interface mode. |
| (config)#interface xe22 | Enter interface mode. |
| (config-if)#ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1 |
| (config-if)#ip address 10.1.1.2/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe20 | Enter interface mode |
| (config-if)#ip address 11.1.1.2/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 2.2.2.2 | Configure OSPF router id same as loopback ip address. |

| | |
|---|--|
| (config-router)#network 2.2.2.2/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 11.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#router-id 2.2.2.2 | Configure the router ID same as loopback IP address. |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure LDP transport address same as loopback address. |
| (config-router)#exit | Exit LDP mode. |
| (config)#interface xe20 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 65000 | Enter BGP router mode. |
| (config-router)#bgp router-id 2.2.2.2 | Configure BGP router-id. |
| (config-router)#neighbor 41.41.41.41 remote-as 65000 | Configure PE2-ASBR1 as an iBGP peer. |
| (config-router)#neighbor 41.41.41.41 update-source lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| (config-router)#address-family vpnv4 | Enter VPNv4 address family mode. |
| (config-router-af)#neighbor 41.41.41.41 activate | Activate the PE neighbor so that it can accept VPN IPv4 routes. |
| (config-router-af)#exit | Exit VPNv4 address family mode. |
| (config-router)#address-family ipv4 vrf vrf1 | Enter the IPv4 address family for VRF vrf1. |
| (config-router-af)#redistribute connected | Redistribute connected route. |
| (config-router-af)#commit | Commit the transaction. |
| (config-router-af)#exit | Exit VPNv4 address family mode. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#commit | Commit candidate configuration to be running configuration. |
| (config)#exit | Exit the config mode. |

P1 Configuration

| | |
|---|--------------------------|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 31.31.31.31/32 secondary | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |

| | |
|--|--|
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#ip address 11.1.1.31/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#ip address 12.1.1.31/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 31.31.31.31 | Configure OSPF router ID same as loopback IP address. |
| (config-router)#network 31.31.31.31/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 11.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 12.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#router-id 31.31.31.31 | Configure the router id same as loopback ip address. |
| (config-router)#transport-address ipv4 31.31.31.31 | Configure LDP transport address same as loopback address. |
| (config-router)#exit | Exit LDP mode. |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-router)#exit | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#commit | Commit the transaction. |
| (config-if)#exit | Exit interface mode. |
| (config)#commit | Commit candidate configuration to be running configuration. |
| (config)#exit | Exit the config mode. |

PE2-ASBR1

| | |
|---|--------------------------|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 41.41.41.41/32 secondary | Assign the IPv4 address. |

| | |
|--|--|
| (config-if)#exit | Exit interface mode. |
| (config)#ip vrf vrf1 | Create a new VRF named vrf1. |
| (config-vrf)#rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| (config-vrf)#route-target both 1:1 | Import routes between route target (RT) ext-communities. |
| (config-vrf)#exit | Exit interface mode. |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1. |
| (config-if)#ip address 20.1.1.41/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#ip address 12.1.1.41/24 | Assign the IPv4 address. |
| (config-if)# | Exit interface mode. |
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 41.41.41.41 | Configure OSPF router ID same as loopback IP address. |
| (config-router)#network 41.41.41.41/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 12.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#transport-address ipv4 41.41.41.41 | Configure LDP transport address same as loopback address. |
| (config-router)#exit | Exit LDP mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-router)#exit | Exit interface mode. |
| (config)#router bgp 65000 | Enter BGP router mode. |
| (config-router)#bgp router-id 41.41.41.41 | Configure BGP router-id. |
| (config-router)#neighbor 2.2.2.2 remote-as 65000 | Configure PE1 as an iBGP peer. |
| (config-router)#neighbor 2.2.2.2 update-source lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| (config-router)#address-family vpnv4 | Enter VPNv4 address family mode. |
| (config-router-af)#neighbor 2.2.2.2 activate | Activate the PE neighbor so that it can accept VPN IPv4 routes. |

| | |
|---|---|
| (config-router-af) #exit | Exit VPNv4 address family mode. |
| (config-router) #address-family ipv4 vrf vrf1 | Enter the IPv4 address family for VRF vrf1. |
| (config-router-af) #neighbor 20.1.1.3 remote-as 65001 | Configure eBGP neighbor. |
| (config-router-af) #redistribute connected | Redistribute connected route. |
| (config-router-af) #commit | Commit the transaction. |
| (config-router-af) #exit-address-family | Exit address family mode. |
| (config-router) #exit | Exit the router mode. |
| (config) #commit | Commit candidate configuration to be running configuration. |
| (config) #exit | Exit the config mode. |

PE3-ASBR2

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config) #interface lo | Enter interface mode. |
| (config-if) #ip address 3.3.3.3/32 secondary | Assign the IPv4 address. |
| (config-if) #exit | Exit interface mode. |
| (config) #ip vrf vrf1 | Create a new VRF named vrf1. |
| (config-vrf) #rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| (config-vrf) #route-target both 1:1 | Import routes between route target (RT) ext-communities. |
| (config-vrf) #exit | Exit interface mode. |
| (config) #interface xe21 | Enter interface mode. |
| (config-if) #ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1. |
| (config-if) #ip address 20.1.1.3/24 | Assign the IPv4 address. |
| (config-if) #exit | Exit interface mode. |
| (config) #interface xe15 | Enter interface mode. |
| (config-if) #ip address 21.1.1.3/24 | Assign the IPv4 address. |
| (config-if) #exit | Exit interface mode |
| (config) #router ospf 1 | Enter router OSPF mode. |
| (config-router) #ospf router-id 3.3.3.3 | Configure OSPF router id same as loopback ip address. |
| (config-router) #network 3.3.3.3/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router) #network 21.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router) #exit | Exit OSPF router mode. |

| | |
|---|---|
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#router-id 3.3.3.3 | Configure the router ID same as loopback IP address. |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure LDP transport address same as loopback address. |
| (config-router)#exit | Exit LDP mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-router)#exit | Exit interface mode. |
| (config)#router rsvp | Enter router RSVP mode. |
| (config-router)#rsvp-trunk lsp1 | Create an RSVP trunk lsp1 and enter the Trunk mode. |
| (config-trunk)#to 5.5.5.5 | Specify the IPv4 egress (destination point-PE4 loopback address) for the LSP. |
| (config-trunk)#exit | Exit interface mode. |
| (config)#router bgp 65001 | Enter BGP router mode. |
| (config-router)#bgp router-id 3.3.3.3 | Configure BGP router-id. |
| (config-router)#neighbor 5.5.5.5 remote-as 65001 | Configure PE4 as an iBGP peer. |
| (config-router)#neighbor 5.5.5.5 update-source lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| (config-router)#address-family vpnv4 | Enter VPNv4 address family mode. |
| (config-router-af)#neighbor 5.5.5.5 activate | Activate the PE neighbor so that it can accept VPN IPv4 routes. |
| (config-router-af)#exit | Exit VPNv4 address family mode. |
| (config-router)#address-family ipv4 vrf vrf1 | Enter the IPv4 address family for VRF vrf1. |
| (config-router-af)#neighbor 20.1.1.41 remote-as 65000 | Configure eBGP neighbor. |
| (config-router-af)#neighbor 20.1.1.41 activate | Activate the eBGP neighbor under address family. |
| (config-router-af)#redistribute connected | Redistribute connected route. |
| (config-router-af)#commit | Commit the transaction. |
| (config-router-af)#exit-address-family | Exit address family mode. |
| (config-router)#exit | Exit the router mode. |
| (config)#commit | Commit candidate configuration to be running configuration. |
| (config)#exit | Exit the config mode. |

P2

| | |
|---------------------|-----------------------|
| #configure terminal | Enter configure mode. |
|---------------------|-----------------------|

| | |
|--|--|
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 40.40.40.40/32 secondary | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#ip address 21.1.1.40/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#ip address 22.1.1.40/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 40.40.40.40 | Configure OSPF router ID same as loopback IP address. |
| (config-router)#network 40.40.40.40/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 21.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 22.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#router-id 40.40.40.40 | Configure the router ID same as loopback IP address. |
| (config-router)#transport-address ipv4 40.40.40.40 | Configure LDP transport address same as loopback address. |
| (config-router)#exit | Exit LDP mode. |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#commit | Commit the transaction. |
| (config-if)#exit | Exit interface mode. |
| (config)#commit | Commit candidate configuration to be running configuration. |
| (config)#exit | Exit the config mode. |

PE4

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 5.5.5.5/32 secondary | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#ip vrf vrf1 | Create a new VRF named vrf1. |
| (config-vrf)#rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| (config-vrf)#route-target both 1:1 | Import routes between route target (RT) ext-communities. |
| (config-vrf)#exit | Exit VRF mode. |
| (config)#interface xe22 | Enter interface mode. |
| (config-if)#ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1. |
| (config-if)#ip address 30.1.1.5/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#ip address 22.1.1.5/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 5.5.5.5 | Configure OSPF router ID same as loopback IP address. |
| (config-router)#network 5.5.5.5/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 22.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit router OSPF mode. |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#router-id 5.5.5.5 | Configure the router ID same as loopback IP address. |
| (config-router)#transport-address ipv4 5.5.5.5 | Configure LDP transport address same as loopback address. |
| (config-router)#exit | Exit LDP mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 65001 | Enter BGP router mode. |
| (config-router)#bgp router-id 5.5.5.5 | Configure BGP router-id. |

| | |
|---|--|
| (config-router)#neighbor 3.3.3.3 remote-as 65001 | Configure PE2-ASBR1 as an iBGP peer. |
| (config-router)#neighbor 3.3.3.3 update-source lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| (config-router)#address-family vpnv4 | Enter VPNv4 address family mode. |
| (config-router-af)#neighbor 3.3.3.3 activate | Activate neighbor. |
| (config-router-af)#exit | Exit VPNv4 Address Family mode. |
| (config-router)#address-family ipv4 vrf vrf1 | Enter IPv4 VRF Address Family mode. |
| (config-router-af)#redistribute connected | Redistribute connected route. |
| (config-router-af)#commit | Commit the transaction. |
| (config-router-af)#exit-address-family | Exit address family mode. |
| (config-router)#exit | Exit the router mode |
| (config)#commit | Commit candidate configuration to be running configuration. |
| (config)#exit | Exit the config mode. |

Validation

PE1

```
#show ip route vrf vrf1 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1"
C    *> 10.1.1.0/24 is directly connected, xe22, 01:05:28
B    *> 20.1.1.0/24 [200/0] via 41.41.41.41, 00:01:18
B    *> 30.1.1.0/24 [200/0] via 41.41.41.41, 00:00:24
C    *> 127.0.0.0/8 is directly connected, lo.vrf1, 01:06:20

Gateway of last resort is not set
#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop         Metric      LocPrf     Weight Path
Route Distinguisher: 1:1 (Default for VRF vrf1)
*> 1 10.1.1.0/24    0.0.0.0           0           100       32768    ?
*>i 20.1.1.0/24    41.41.41.41        0           100         0        ?
*>i 30.1.1.0/24    41.41.41.41        0           100         0       65001 ?
  Announced routes count = 1
  Accepted routes count = 2
Route Distinguisher: 1:1
*>i 20.1.1.0/24    41.41.41.41        0           100         0        ?
*>i 30.1.1.0/24    41.41.41.41        0           100         0       65001 ?
  Announced routes count = 0
  Accepted routes count = 2
```

PE2-ASBR1

```
#show ip route vrf vrf1 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1"
C   *> 20.1.1.0/24 is directly connected, xe22, 01:05:28
B   *> 10.1.1.0/24 [200/0] via 2.2.2.2, 00:01:18
B   *> 30.1.1.0/24 [20/0] via 20.1.1.3, xe2, 00:54:13
C   *> 127.0.0.0/8 is directly connected, lo.vrf1, 01:06:20

Gateway of last resort is not set

#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

      Network          Next Hop           Metric    LocPrf    Weight Path
Route Distinguisher: 1:1 (Default for VRF vrf1)
*>i 10.1.1.0/24        2.2.2.2             0          100        0        ?
*> 1 20.1.1.0/24       0.0.0.0             0          100       32768        ?
*   20.1.1.0/24       20.1.1.3            0          100        0       65001 ?
*> 1 30.1.1.0/24       20.1.1.3            0          100        0       65001 ?
Announced routes count = 3
Accepted routes count = 1
Route Distinguisher: 1:1
*>i 10.1.1.0/24        2.2.2.2             0          100        0        ?
Announced routes count = 0
Accepted routes count = 1
```

PE3-ASBR2

```
#show ip route vrf vrf1 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1"
B   *> 10.1.1.0/24 [20/0] via 20.1.1.41, xe22, 00:55:54
C   *> 20.1.1.0/24 is directly connected, xe22, 01:05:28
B   *> 30.1.1.0/24 [200/0] via 5.5.5.5, 00:01:18
C   *> 127.0.0.0/8 is directly connected, lo.vrf1, 01:06:20

Gateway of last resort is not set

#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

      Network          Next Hop           Metric    LocPrf    Weight Path
Route Distinguisher: 1:1 (Default for VRF vrf1)
*> 1 10.1.1.0/24       20.1.1.41           0          100        0       65000 ?
*> 1 20.1.1.0/24       0.0.0.0             0          100       32768        ?
```

```
*      20.1.1.0/24      20.1.1.41      0      100      0      65000 ?
*>i  30.1.1.0/24      5.5.5.5      0      100      0      ?
Announced routes count = 3
Accepted routes count = 1
Route Distinguisher: 1:1
*>i  30.1.1.0/24      5.5.5.5      0      100      0      ?
Announced routes count = 0
Accepted routes count = 1
```

PE4

```
#show ip route vrf vrf1 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1"
B      *>  10.1.1.0/24 [200/0] via 3.3.3.3, 00:00:08
B      *>  20.1.1.0/24 [200/0] via 3.3.3.3, 00:02:45
C      *>  30.1.1.0/24 is directly connected, xe18, 01:02:20
C      *>  127.0.0.0/8 is directly connected, lo.vrf1, 01:05:36

Gateway of last resort is not set

#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

      Network          Next Hop          Metric      LocPrf      Weight Path
Route Distinguisher: 1:1 (Default for VRF vrf1)
*>i  10.1.1.0/24      3.3.3.3          0          100          0      65000 ?
*>i  20.1.1.0/24      3.3.3.3          0          100          0      ?
*> 1 30.1.1.0/24      0.0.0.0          0          100          32768  ?
Announced routes count = 1
Accepted routes count = 2
Route Distinguisher: 1:1
*>i  10.1.1.0/24      3.3.3.3          0          100          0      65000 ?
*>i  20.1.1.0/24      3.3.3.3          0          100          0      ?
Announced routes count = 0
Accepted routes count = 2
```

Inter-AS VPN Option-B Overview

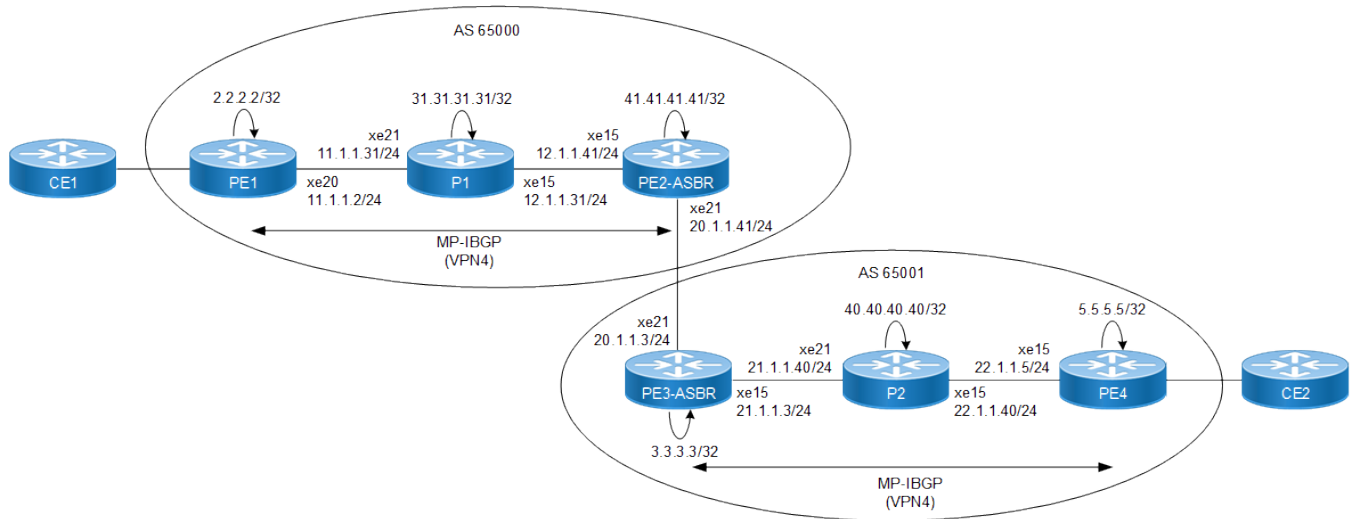
- Inter-AS Option B is a more scalable solution compared to Option A. It does not require any VRFs on the ASBRs, it uses VPNv4 eBGP to exchange VPNv4 updates.
- Single interface to connect the ASBRs.
- Packets are sent labelled between the ASBRs.
- No need for VRFs on the ASBR.
- ASBRs must be directly connected.
- Scales better than Option A.



Note: ASBR BGP label allocation mode uses per-prefix label allocation, and per-VRF label allocation cannot be achieved in ASBRs (VPNv4 and VPNv6).

Topology


Figure 90. InterAS-VPN Option-B



Inter-AS VPN Option-B Configuration

PE1

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 2.2.2.2/32 secondary | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#ip vrf vrf1 | Create a new VRF named vrf1. |
| (config-vrf)#rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| (config-vrf)#route-target both 1:1 | Import routes between route target (RT) ext-communities. |
| (config-vrf)#exit | Exit VRF mode. |
| (config)#interface xe22 | Enter interface mode. |
| (config-if)#ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1. |
| (config-if)#ip address 10.1.1.2/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe20 | Enter interface mode. |
| (config-if)#ip address 11.1.1.2/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |

| | |
|--|--|
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 2.2.2.2 | Configure OSPF router id same as loopback ip address. |
| (config-router)#network 2.2.2.2/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 11.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| |  Note: It is preferable to use the loopback address as the transport address |
| (config-router)#targeted-peer ipv4 41.41.41.41 | Configure targeted peer. |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configure targeted peer. |
| (config-router-targeted-peer)#targeted-peer ipv4 5.5.5.5 | Configure targeted peer. |
| (config-router-targeted-peer)#targeted-peer ipv4 41.41.41.41 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode. |
| (config-router)#exit | Exit router mode |
| (config)#interface xe20 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 65000 | Enter BGP router mode. |
| (config-router)#bgp router-id 2.2.2.2 | Configure BGP router-id. |
| (config-router)#neighbor 41.41.41.41 remote-as 65000 | Configure PE2-ASBR1 as an iBGP peer. |
| (config-router)#neighbor 41.41.41.41 update-source lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| (config-router)#address-family vpnv4 | Enter VPNv4 address family mode. |
| (config-router-af)#neighbor 41.41.41.41 activate | Activate the PE neighbor so that it can accept VPN IPv4 routes. |
| (config-router-af)#exit | Exit VPNv4 address family mode. |
| (config-router)#address-family ipv4 vrf vrf1 | Enter the IPv4 address family for VRF vrf1. |


| | |
|--|---|
| (config-router-af) #redistribute connected | Redistribute connected route. |
| (config-router-af) #commit | Commit the transaction. |
| (config-router-af) #exit-address-family | Exit address family mode. |
| (config-router) #exit | Exit the router mode. |
| (config) #commit | Commit candidate configuration to be running configuration. |
| (config) #exit | Exit the config mode. |

P1

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| (config) #interface lo | Enter interface mode. |
| (config-if) #ip address 31.31.31.31/32 secondary | Assign the IPv4 address. |
| (config-if) #exit | Exit interface mode. |
| (config) #interface xe21 | Enter interface mode. |
| (config-if) #ip address 11.1.1.31/24 | Assign the IPv4 address. |
| (config-if) #exit | Exit interface mode. |
| (config) #interface xe15 | Enter interface mode. |
| (config-if) #ip address 12.1.1.31/24 | Assign the IPv4 address. |
| (config-if) #exit | Exit interface mode. |
| (config) #router ospf 1 | Enter router OSPF mode. |
| (config-router) #ospf router-id 31.31.31.31 | Configure OSPF router id same as loopback ip address. |
| (config-router) #network 31.31.31.31/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router) #network 11.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router) #network 12.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router) #exit | Exit OSPF router mode. |
| (config) #router ldp | Enter router LDP mode. |
| (config-router) #router-id 31.31.31.31 | Set the router ID to IP address 31.31.31.31 |
| (config-router) #transport-address ipv4 31.31.31.31 | Configure LDP transport address same as loopback address. |
| (config-router) #exit | Exit LDP mode. |
| (config) #interface xe21 | Enter interface mode. |
| (config-if) #label-switching | Enable label switching in interface. |
| (config-if) #enable-ldp ipv4 | Enable LDP in interface. |
| (config-if) #exit | Exit interface mode. |

| | |
|-----------------------------|---|
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#commit | Commit the transaction. |
| (config-if)#exit | Exit interface mode. |
| (config)#commit | Commit candidate configuration to be running configuration. |
| (config)#exit | Exit the config mode. |

PE2-ASBR

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 41.41.41.41/32 secondary | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 41.41.41.41 | Set the router ID to IP address 41.41.41.41 |
| (config-router)#transport-address ipv4 41.41.41.41 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Configure targeted peer. |
| (config-router-targeted-peer)#targeted-peer ipv4 3.3.3.3 | Configure targeted peer. |
| (config-router-targeted-peer)#targeted-peer ipv4 5.5.5.5 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#ip address 12.1.1.41/24 | Assign ipv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 41.41.41.41 | Configure OSPF router id same as loopback ip address. |
| (config-router)#network 41.41.41.41/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 12.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |

| | |
|---|--|
| (config-router)#exit | Exit OSPF router mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#ip address 20.1.1.41/24 | Assign the IPv4 address. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 65000 | Enter BGP router mode. |
| (config-router)#bgp router-id 41.41.41.41 | Configure BGP router-id. |
| (config-router)#no bgp inbound-route-filter | Disable inbound route filter. |
| (config-router)#neighbor 2.2.2.2 remote-as 65000 | Configure PE1 as an iBGP peer. |
| (config-router)#neighbor 2.2.2.2 update-source lo | Update the source as loopback for iBGP peering with the remote PE1 router. |
| (config-router)#neighbor 20.1.1.3 remote-as 65001 | Configure eBGP neighbor with ASBR2. |
| (config-router)#address-family vpnv4 | Enter VPNv4 address family mode. |
| (config-router-af)#neighbor 2.2.2.2 activate | Activate the PE neighbor so that it can accept VPN IPv4 routes. |
| (config-router-af)#neighbor 2.2.2.2 next-hop-self | Configure this to make the router the next hop for a BGP neighbor. |
| (config-router-af)#neighbor 20.1.1.3 activate | Activate the ASBR eBGP neighbor. |
| (config-router-af)#neighbor 20.1.1.3 allow-ebgp-vpn | Configure this to allow exchange of VPN updates between eBGP peers. |
| (config-router-af)#neighbor 20.1.1.3 activate | Activate the ASBR eBGP neighbor. |
| (config-router-af)#commit | Commit candidate configuration to the running configuration. |


PE3-ASBR

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 3.3.3.3/32 secondary | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 3.3.3.3 | Set the router ID to IP address 3.3.3.3 |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |

| | |
|--|---|
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Configure targeted peer. |
| (config-router-targeted-peer)#targeted-peer ipv4 41.41.41.41 | Configure targeted peer. |
| (config-router-targeted-peer)#targeted-peer ipv4 5.5.5.5 | |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#ip address 20.1.1.3/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#ip address 21.1.1.3/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 3.3.3.3 | Configure OSPF router id same as loopback ip address. |
| (config-router)#network 3.3.3.3/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 21.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#router ldp | Enter router ldp mode. |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure LDP transport address same as loopback address |
| (config-router)#exit | Exit LDP mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-router)#exit | Exit LDP mode. |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 65001 | Enter BGP router mode. |
| (config-router)#bgp router-id 3.3.3.3 | Configure BGP router-id. |
| (config-router)#no bgp inbound-route-filter | Disable inbound route filter. |

| | |
|--|--|
| (config-router)#neighbor 5.5.5.5 remote-as 65001 | Configure PE4 as an iBGP peer. |
| (config-router)#neighbor 5.5.5.5 update-source lo | Update the source as loopback for iBGP peering with the remote PE1 router. |
| (config-router)#neighbor 20.1.1.41 remote-as 65000 | Configure eBGP neighbor with PE4. |
| (config-router)#address-family vpnv4 | Enter VPNv4 address family mode. |
| (config-router-af)#neighbor 5.5.5.5 activate | Activate the PE neighbor so that it can accept VPN IPv4 routes. |
| (config-router-af)#neighbor 5.5.5.5 next-hop-self | Configure this to make the router the next hop for a BGP neighbor. |
| (config-router-af)#neighbor 20.1.1.41 activate | Activate the ASBR eBGP neighbor. |
| (config-router-af)#neighbor 20.1.1.41 allow-ebgp-vpn | Configure this to allow exchange of vpn updates between eBGP peers. |
| (config-router-af)#neighbor 20.1.1.41 activate | Activate the ASBR eBGP neighbor. |
| (config-router-af)#commit | Commit candidate configuration to be running configuration. |


P2

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 40.40.40.40/32 secondary | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 40.40.40.40 | Set the router ID to IP address 40.40.40.40 |
| (config-router)#transport-address ipv4 40.40.40.40 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |
| <div>  Note: It is preferable to use the loopback address as the transport address. </div> | |
| (config-router)#targeted-peer ipv4 40.40.40.40 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#ip address 21.1.1.40/24 | Assign the IPv4 address. |
| (config-if)#exit (config-router)#targeted-peer ipv4 40.40.40.40 | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |

| | |
|--|--|
| (config-if)#ip address 22.1.1.40/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 40.40.40.40 | Configure OSPF router id same as loopback ip address. |
| (config-router)#network 40.40.40.40/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 21.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 22.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#transport-address ipv4 40.40.40.40 | Configure LDP transport address same as loopback address. |
| (config-router)#exit | Exit LDP mode. |
| (config)#interface xe21 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable ldp in interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#commit | Commit the transaction. |
| (config-if)#exit | Exit interface mode. |
| (config)#commit | Commit candidate configuration to be running configuration. |
| (config)#exit | Exit the config mode. |

PE4

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 5.5.5.5/32 secondary | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 5.5.5.5 | Set the router ID to IP address 5.5.5.5 |
| (config-router)#transport-address ipv4 5.5.5.5 | Configure the transport address for IPV4 (for IPV6 use ipv6) to be used for a TCP session over which LDP will run. |

| | |
|---|---|
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#targeted-peer ipv4 5.5.5.5 | Configure targeted peer. |
| (config-router-targeted-peer)# targeted-peer ipv4 2.2.2.2 | Configure targeted peer. |
| (config-router-targeted-peer)# targeted-peer ipv4 41.41.41.41 | Configure targeted peer. |
| (config-router-targeted-peer)# targeted-peer ipv4 3.3.3.3 | Configure targeted peer. |
| (config-router)#exit | Exit router mode |
| (config)#ip vrf vrf1 | Create a new VRF named vrf1. |
| (config-vrf)#rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| (config-vrf)#route-target both 1:1 | Import routes between route target (RT) ext-communities. |
| (config-vrf)#exit | Exit VRF mode. |
| (config)#interface xe22 | Enter interface mode. |
| (config-if)#ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1. |
| (config-if)#ip address 30.1.1.5/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#ip address 22.1.1.5/24 | Assign the IPv4 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Enter router OSPF mode. |
| (config-router)#ospf router-id 5.5.5.5 | Configure OSPF router id same as loopback ip address. |
| (config-router)#network 5.5.5.5/32 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#network 22.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 65001 | Enter BGP router mode. |
| (config-router)#bgp router-id 5.5.5.5 | Configure BGP router-id. |

| | |
|---|--|
| (config-router)#neighbor 3.3.3.3 remote-as 65001 | Configure PE2-ASBR1 as an iBGP peer. |
| (config-router)#neighbor 3.3.3.3 update-source lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| (config-router)#address-family vpnv4 | Enter VPNv4 address family mode. |
| (config-router-af)#neighbor 3.3.3.3 activate | Activate the PE neighbor so that it can accept VPN IPv4 routes. |
| (config-router-af)#exit | Exit VPNv4 address family mode. |
| (config-router)#address-family ipv4 vrf vrf1 | Enter the IPv4 address family for VRF vrf1. |
| (config-router-af)#redistribute connected | Redistribute connected route. |
| (config-router-af)#commit | Commit the transaction. |
| (config-router-af)#exit-address-family | Exit address family mode. |
| (config-router)#exit | Exit the router mode. |
| (config)#commit | Commit candidate configuration to be running configuration. |
| (config)#exit | Exit the config mode. |

Validation

PE1

```

PE1#sh ldp session
Peer IP Address   IF Name   My Role   State       KeepAlive   UpTime
31.31.31.31      xe20      Passive   OPERATIONAL 30          00:49:02

#show ip route vrf vrf1 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1"
C    *> 10.1.1.0/24 is directly connected, xe22, 03:49:26
B    *> 30.1.1.0/24 [200/0] via 41.41.41.41, 00:00:41
C    *> 127.0.0.0/8 is directly connected, lo.vrf1, 03:50:18

PE1#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

      Network                Next Hop           Metric    LocPrf   Weight Path
Route Distinguisher: 1:1 (Default for VRF vrf1)
*> 1 10.1.1.0/24              0.0.0.0             0         100     32768    ?
*>i 30.1.1.0/24              41.41.41.41         0         100      0      65001 ?
  Announced routes count = 1
  Accepted routes count = 1
Route Distinguisher: 1:1
*>i 30.1.1.0/24              41.41.41.41         0         100      0      65001 ?
  Announced routes count = 0
  Accepted routes count = 1

```


PE2-ASBR1

```

PE2-ASBR2#show ldp session
Peer IP Address  IF Name  My Role  State  KeepAlive  UpTime
31.31.31.31      xe47    Active   OPERATIONAL  30        00:16:22
3.3.3.3          xe45    Active   OPERATIONAL  30        00:14:54

#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

      Network          Next Hop          Metric    LocPrf    Weight Path
Route Distinguisher: 1:1
*>i 10.1.1.0/24        2.2.2.2           0         100       0         ?
*> 1 30.1.1.0/24       20.1.1.3          0         100       0        65001 ?
Announced routes count = 0
Accepted routes count = 2

#show ip bgp vpnv4 all summary
BGP router identifier 41.41.41.41, local AS number 65000
BGP table version is 4
2 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv  MsgSen  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
2.2.2.2           4 65000  168    171     4       0    0    00:29:03      1
20.1.1.3          4 65001  111    119     4       0    0    00:42:51      1

Total number of neighbors 2

Total number of Established sessions 2

```

PE3-ASBR2

```

PE3-ASBR2#sh ldp session
Peer IP Address  IF Name  My Role  State  KeepAlive  UpTime
41.41.41.41      xe45    Passive  OPERATIONAL  30        00:15:47
40.40.40.40      xe6     Passive  OPERATIONAL  30        00:24:32

#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

      Network          Next Hop          Metric    LocPrf    Weight Path
Route Distinguisher: 1:1
*> 1 10.1.1.0/24        20.1.1.41         0         100       0        65000 ?
*>i 1 30.1.1.0/24       5.5.5.5           0         100       0         ?
Announced routes count = 0
Accepted routes count = 2

#show ip bgp vpnv4 all summary
BGP router identifier 3.3.3.3, local AS number 65001
BGP table version is 4
2 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv  MsgSen  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
5.5.5.5           4 65001   41     45     4       0    0    00:15:59      1
20.1.1.41         4 65000  115    118     4       0    0    00:43:58      1

Total number of neighbors 2

```

Total number of Established sessions 2

PE4

```
#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop        Metric      LocPrf     Weight Path
Route Distinguisher: 1:1 (Default for VRF vrf1)
*>i 10.1.1.0/24      3.3.3.3          0           100         0        65000 ?
*> 1 30.1.1.0/24    0.0.0.0          0           100        32768      ?
Announced routes count = 1
Accepted routes count = 1
Route Distinguisher: 1:1
*>i 10.1.1.0/24      3.3.3.3          0           100         0        65000 ?
Announced routes count = 0
Accepted routes count = 1

#show ip route vrf vrf1 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1"
B   *> 10.1.1.0/24 [200/0] via 3.3.3.3, 00:00:48
C   *> 30.1.1.0/24 is directly connected, xe22, 03:46:38
C   *> 127.0.0.0/8 is directly connected, lo.vrf1, 03:49:54

Gateway of last resort is not set

PE4#sh ldp session

Peer IP Address  IF Name    My Role    State        KeepAlive    UpTime
40.40.40.40      xe5        Passive    OPERATIONAL  30           00:26:30

P#sh ldp session
Peer IP Address  IF Name    My Role    State        KeepAlive    UpTime
2.2.2.2          xe7        Active     OPERATIONAL  30           00:31:27
41.41.41.41      xe1        Passive    OPERATIONAL  30           00:15:12

P#
P2#sh ldp session
Peer IP Address  IF Name    My Role    State        KeepAlive    UpTime
3.3.3.3          xe6        Active     OPERATIONAL  30           00:24:58
5.5.5.5          xe5        Active     OPERATIONAL  30           00:24:15
```

Inter-AS VPN Option-C Configuration

This is an option which enables BGP VPNv4 routes exchange between the two or more Provider Edge routers residing in different AS'es. There is no VPNv4 route exchange between the ASBRs and hence the ASBRs are completely transparent of the VPNs provisioned on the connecting AS'es. This option is more scalable than the previous two options because it removes the restriction of installing VPNv4 LFIB entries on the ASBRs hence reducing the resource requirements on them. ASBRs are only required to forward data based on the transport LSPs configured.

Pre-requisites to successfully configure this option.

- There should be end to end transport LSP between the two PE's residing in different AS'es.
- Could use BGP Labeled-unicast to exchange labels for the two PE's residing in different AS'es.
- This option can utilize RR's to exchange VPNv4 routes between the two AS'es and hence should have capability to keep the next-hops unchanged when reflecting routes from iBGP to eBGP
- LDP/RSVP/SR can be used as a transport within the AS.

There are multiple ways in which this option can be configured.

- Inter-AS Option-C with VPNv4 Routes exchange between the two RRs in different AS'es
- Inter-AS Option-C with VPNv4 Routes exchange between two PE's in different AS.
- Inter-AS Option-C with VPNv4 Routes exchange between PE of one AS and RR of another AS.

This section explains the configuration for the first two points.

Inter-AS Option-C with RR Configuration

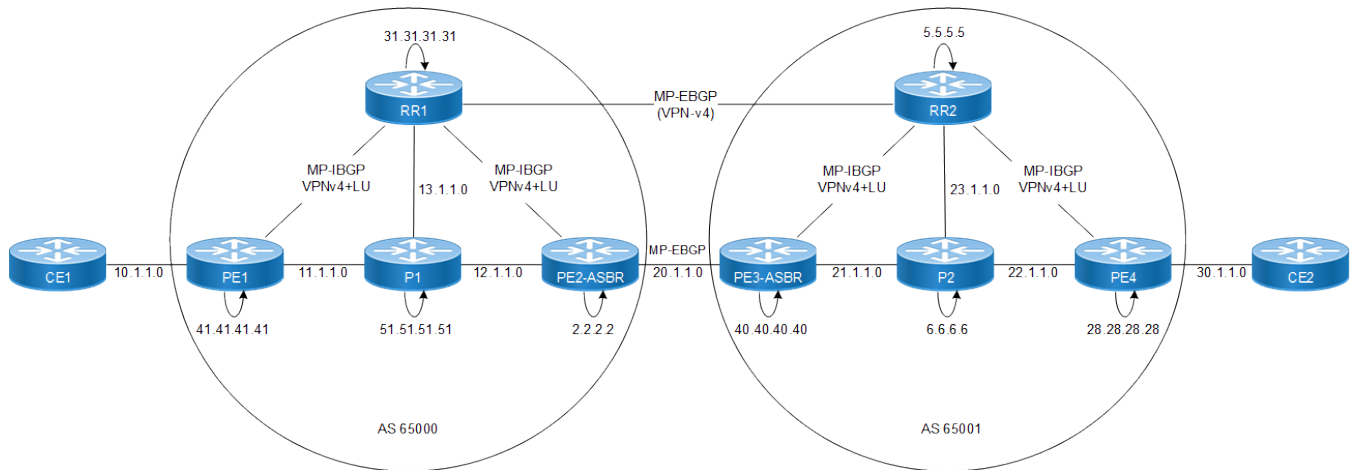
This sub-option enables VPNv4 routes exchange between the Route-Reflectors of different AS'es and thus requires

1. Remote PE loopback addresses to be known in the AS - This is accomplished using BGP-LU and is advertised inside the AS using the Route-reflector and outside the AS using the eBGP Peering between the ASBRs.
2. MP-eBGP peering between the two RRs of different AS'es in order to exchange VPNv4 routes.
3. To be able to have an eBGP session between RR's, Remote Route-reflector Loopback addresses are required to be known in the local AS and vice-versa. This is accomplished using the mutual redistribution between OSPF and bgp on the two ASBRs. The redistribution is controlled using the route-maps and only permits the
 - a. Local Route-reflectors Loopback address to be redistributed from local OSPF to BGP.
 - b. Remote Route-reflectors Loopback address to be redistributed from bgp to local OSPF.
4. Next-hop attribute should not be changed when RR reflects the VPNv4 route from iBGP peering to the remote RR over an eBGP peering. This is accomplished using the configuration "attribute-unchanged next-hop" for the eBGP neighbor.
5. Next-hop attribute should be changed when ASBR advertises the remote PE's Loopback address as LU route to the RR. This is accomplished using the configuration "next-hop self" for the iBGP RR neighbor

The first and fifth points enables an end to end LSP between the two PE's which needs to run the L3VPN service and (2), (3) and (4) enables the RR's to form BGP neighborhood with each other and successfully exchange the VPNv4 Routes.

Topology

Figure 91. InterAS-VPN Option-C with RR



PE1

| | |
|---|--|
| PE1(#configure terminal | Enter configure mode. |
| PE1(config-router)#interface lo | Enter interface mode. |
| PE1(config-if)# ip address 41.41.41.41/32 secondary | Assign the IPv4 address. |
| PE1(config-if)#exit | Exit interface mode. |
| PE1(config)#ip vrf vrf1 | Create a new VRF named vrf1. |
| PE1(config-vrf)# rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| PE1(config-vrf)# route-target both 1:1 | Configure import and export route-target values |
| PE1(config-vrf)#exit | Exit VRF mode. |
| PE1(config-if)#interface xe23 | Enter interface mode. |
| PE1(config-if)# ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1. |
| PE1(config-if)# ip address 10.1.1.2/24 | Assign the IPv4 address. |
| PE1(config-if)#exit | Exit interface mode. |
| PE1(config-if)#interface xe15 | Enter interface mode. |
| PE1(config-if)# ip address 11.1.1.2/24 | Assign the IPv4 address. |
| PE1(config-if)#exit | Exit interface mode. |
| PE1(config-if)#router ospf 1 | Enter router OSPF mode. |
| PE1(config-router)# ospf router-id 41.41.41.41 | Configure OSPF router id same as loopback ip address. |
| PE1(config-router)# network 11.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE1(config-router)# network 41.41.41.41/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |

| | |
|---|--|
| PE1(config)#router ldp | Enter router LDP mode. |
| PE1(config-router)#router-id 41.41.41.41 | Set the router ID to IP address 41.41.41.41 |
| PE1(config-router)#transport-address ipv4 41.41.41.41 | Configure LDP transport address same as loopback address. |
| PE1(config-router)#exit | Exit LDP mode. |
| PE1(config-if)#interface xe15 | Enter interface mode. |
| PE1(config-if)# label-switching | Enable label switching in interface. |
| PE1(config-if)# enable-ldp ipv4 | Enable LDP in interface. |
| PE1(config-if)#exit | Exit interface mode. |
| PE1(config-router)#router bgp 65000 | Enter BGP router mode. |
| PE1(config-router)# bgp router-id 41.41.41.41 | Configure BGP router-id. |
| PE1(config-router)# no bgp default ipv4-unicast | Configure BGP peer to have no ipv4 unicast capability exchange by default. |
| PE1(config-router)# neighbor 31.31.31.31 remote-as 65000 | Configure neighbor with remote AS. |
| PE1(config-router)# neighbor 31.31.31.31 update-source lo | Update the loopback as the source of BGP session. |
| PE1(config-router)# allocate-label all | Allocate the label for all advertised networks. |
| PE1(config-router)# address-family ipv4 labeled-unicast | Enter the IPv4 labeled-unicast address family |
| PE1(config-router-af)# neighbor 31.31.31.31 activate | Activate the neighbor under the address-family. |
| PE1(config-router-af)# exit-address-family | Exit IPv4 labeled-unicast Address Family mode. |
| PE1(config-router)# address-family vpnv4 unicast | Enter vpnv4 unicast address-family mode. |
| PE1(config-router-af)# neighbor 31.31.31.31 activate | Activate the neighbor under vpnv4 unicast address-family. |
| PE1(config-router-af)# exit-address-family | Exit vpnv4 unicast address-family. |
| PE1(config-router)# address-family ipv4 vrf vrf1 | Enter ipv4 VRF address-family mode. |
| PE1(config-router-af)# redistribute connected | Redistribute connected networks under VRF address-family. |
| PE1(config-router-af)# exit-address-family | Exit ipv4 VRF address-family. |
| PE1(config-router)#exit | Exit the router mode. |
| PE1(config)#commit | Commit candidate configuration to be running configuration. |
| PE1(config)#exit | Exit the config mode. |

P1

| | |
|--|--|
| P1(config)#interface lo | Enter interface mode. |
| P1(config-if)# ip address 51.51.51.51/32 secondary | Assign the secondary IPv4 address on loopback interface. |

| | |
|--|---|
| P1(config-if)#interface xe15 | Enter interface mode. |
| P1(config-if)# ip address 11.1.1.31/24 | Assign the IPv4 address. |
| P1(config-if)#interface xe21 | Enter interface mode. |
| P1(config-if)# ip address 12.1.1.31/24 | Assign the IPv4 address. |
| P1(config-if)#interface xe22 | Enter interface mode. |
| P1(config-if)# ip address 13.1.1.1/24 | Assign the IPv4 address. |
| P1(config-if)#router ospf 1 | Enter router OSPF mode. |
| P1(config-router)# ospf router-id 51.51.51.51 | Configure OSPF router-id (optional). |
| P1(config-router)# network 11.1.1.0/24 area 0.0.0.0 | Enable networks for OSPF protocol. |
| P1(config-router)# network 12.1.1.0/24 area 0.0.0.0 | Enable networks for OSPF protocol. |
| P1(config-router)# network 13.1.1.0/24 area 0.0.0.0 | Enable networks for OSPF protocol. |
| P1(config-router)# network 51.51.51.51/32 area 0.0.0.0 | Enable networks for OSPF protocol. |
| P1(config-router)#router ldp | Enter router LDP mode. This is required to enable LDP globally on the router. |
| P1(config-router)#router-id 51.51.51.51 | Set the router ID to IP address 51.51.51.51 |
| P1(config-router)#transport-address ipv4 51.51.51.51 | Configure LDP transport address same as loopback address. |
| P1(config-router)#interface xe15 | Enter interface configuration mode. |
| P1(config-if)# label-switching | Configure label-switching on the interface. |
| P1(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| P1(config-if)#interface xe21 | Enter interface configuration mode. |
| P1(config-if)# label-switching | Configure label-switching on the interface. |
| P1(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| P1(config-if)#exit | Exit interface mode. |
| P1(config)#commit | Commit candidate configuration to the running configuration. |
| P1(config)#exit | Exit the config mode. |

RR1

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| RR1(config)#interface lo | Enter interface mode. |
| RR1(config-if)# ip address 31.31.31.31/32 secondary | Assign the IPv4 address. |
| RR1(config-if)#interface xe22 | Exit interface mode. |
| RR1(config-if)# ip address 13.1.1.2/24 | Assign the IPv4 address. |
| RR1(config-if)#router ospf 1 | Enter router OSPF mode. |
| RR1(config-router)# ospf router-id 31.31.31.31 | Configure OSPF router-id explicitly (optional). |
| RR1(config-router)# network 13.1.1.0/24 area 0 | Enable networks for OSPF protocol under area 0. |

| | |
|--|--|
| 0.0.0.0 | |
| RR1(config-router)# network 31.31.31.31/32 area 0.0.0.0 | Enable networks for OSPF protocol under area 0. |
| RR1(config-router)#router bgp 65000 | Enter BGP router mode. |
| RR1(config-router)# bgp router-id 31.31.31.31 | Configure BGP router-id. |
| RR1(config-router)# no bgp default ipv4-unicast | Disable default ipv4-unicast capability exchange with BGP neighbors. |
| RR1(config-router)# no bgp inbound-route-filter | Disable inbound route-filtering for bgp VPNV4 routes. This is required for Route-reflectors. |
| RR1(config-router)# neighbor 2.2.2.2 remote-as 65000 | Configure PE2-ASBR as iBGP neighbor. |
| RR1(config-router)# neighbor 2.2.2.2 update-source lo | Configure BGP speaker to use loopback address as source for BGP session. |
| RR1(config-router)# neighbor 5.5.5.5 remote-as 65001 | Configure RR2 as an eBGP neighbor. |
| RR1(config-router)# neighbor 5.5.5.5 ebgp-multihop | Enable multihop on eBGP session. |
| RR1(config-router)# neighbor 5.5.5.5 update-source lo | Configure BGP speaker to use loopback address as source for this BGP session. |
| RR1(config-router)# neighbor 41.41.41.41 remote-as 65000 | Configure PE1 as an iBGP neighbor. |
| RR1(config-router)# neighbor 41.41.41.41 update-source lo | Configure BGP speaker to use loopback address as source for this BGP session. |
| RR1(config-router)# allocate-label all | Configure this peer to allocate label for all advertised networks. |
| RR1(config-router)# address-family ipv4 labeled-unicast | Enter ipv4 labeled-unicast address-family. |
| RR1(config-router-af)# neighbor 2.2.2.2 activate | Activate PE2-ASBR for this AFI/SAFI. |
| RR1(config-router-af)# neighbor 2.2.2.2 route-reflector-client | Configure PE2-ASBR as Route-reflector client. |
| RR1(config-router-af)# neighbor 41.41.41.41 activate | Activate PE1 for this AFI/SAFI. |
| RR1(config-router-af)# neighbor 41.41.41.41 route-reflector-client | Configure PE1 as Route-reflector client. |
| RR1(config-router-af)# exit-address-family | Exit ipv4 labeled-unicast address-family. |
| RR1(config-router)# address-family vpnv4 unicast | Enter vpnv4 unicast address-family. |
| RR1(config-router-af)# neighbor 5.5.5.5 allow-ebgp-vpn | Allow eBGP VPN route exchange for the RR2 eBGP neighbor. |
| RR1(config-router-af)# neighbor 5.5.5.5 activate | Activate the RR2 eBGP neighbor. |
| RR1(config-router-af)# neighbor 5.5.5.5 attribute-unchanged next-hop | Configure next-hop unchanged for this eBGP neighbor for all vpnv4 NLRIs to keep original next-hop while advertising to this eBGP neighbor. |
| RR1(config-router-af)# neighbor 41.41.41.41 activate | Activate PE1 iBGP neighbor. |

| | |
|--|---|
| RR1(config-router-af)# neighbor 41.41.41.41 route-reflector-client | Configure PE1 as Route-reflector client for this AFI/SAFI. |
| RR1(config-router-af)# exit-address-family | Exit Address-family mode. |
| RR1(config-router)#exit | Exit the router mode. |
| RR1(config)#commit | Commit candidate configuration to be running configuration. |
| RR1(config)#exit | Exit the config mode. |

PE2-ASBR

| | |
|--|--|
| PE2-ASBR#configure terminal | Enter configure mode. |
| PE2-ASBR(config)#interface lo | Enter interface mode. |
| PE2-ASBR(config-if)# ip address 2.2.2.2/32 secondary | Assign a secondary IPv4 address to the loopback interface. |
| PE2-ASBR(config-if)#interface xe20 | Enter interface mode. |
| PE2-ASBR(config-if)# ip address 12.1.1.2/24 | Assign an IPv4 address to the interface. |
| PE2-ASBR(config-if)#interface xe5 | Enter interface mode. |
| PE2-ASBR(config-if)# ip address 20.1.1.2/24 | Assign an IPv4 address to the interface. |
| PE2-ASBR(config-if)#ip access-list standard RR1_LO | Create an IP standard access list to permit RR1s loopback address. This is required to redistribute RR1 Lo from OSPF to BGP. |
| PE2-ASBR(config-ip-acl-std)# permit host 31.31.31.31 | Permit RR1 loopback address. |
| PE2-ASBR(config-ip-acl-std)#ip access-list standard RR2_LO | Create an IP standard access list to permit RR2 Lo address. This is required to redistribute RR2 Lo from BGP to OSPF. |
| PE2-ASBR(config-ip-acl-std)# permit host 5.5.5.5 | Permit RR2s Loopback address. |
| PE2-ASBR(config-ip-acl-std)#route-map RR1_LO permit 10 | Create a Route-map to allow RR1-LO address . This is required when redistributing from OSPF to BGP. |
| PE2-ASBR(config-route-map)# match ip address RR1_LO | match the RR1_LO access list. |
| PE2-ASBR(config-route-map)#route-map RR2_LO permit 10 | Create a Route-map to allow RR2-LO address . This is required when redistributing from BGP to OSPF. |
| PE2-ASBR(config-route-map)# match ip address RR2_LO | match the RR2_LO access list. |
| PE2-ASBR(config-route-map)#router ospf 1 | Enter router OSPF mode. |
| PE2-ASBR(config-router)# ospf router-id 2.2.2.2 | Configure OSPF router-id explicitly (optional). |
| PE2-ASBR(config-router)# redistribute bgp route-map RR2_LO | Redistribute RR2's Lo from BGP to OSPF. This is required for eBGP session between RR1 and RR2. |
| PE2-ASBR(config-router)# network 2.2.2.2/32 area 0.0.0.0 | Enable OSPF networks/links in area 0. |
| PE2-ASBR(config-router)# network 12.1.1.0/24 area 0.0.0.0 | Enable OSPF networks/links in area 0. |
| PE2-ASBR(config-router)#router ldp | Enter Router ldp mode. This is required to enable ldp |

| | |
|--|---|
| | globally on the router. |
| PE2-ASBR(config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| PE2-ASBR(config-router)#transport-address ipv4 2.2.2.2 | Configure LDP transport address same as loopback address. |
| PE2-ASBR(config-router)#interface xe20 | Enter interface mode. |
| PE2-ASBR(config-if)# label-switching | Configure label-switching to enable MPLS forwarding on this interface. |
| PE2-ASBR(config-if)# enable-ldp ipv4 | Enable LDP on this interface. |
| PE2-ASBR(config-if)#interface xe5 | Enter interface mode. |
| PE2-ASBR(config-if)# label-switching | Configure label-switching to enable MPLS forwarding on this interface. |
| PE2-ASBR(config-if)#router bgp 65000 | Enter router BGP configuration mode. |
| PE2-ASBR(config-router)# bgp router-id 2.2.2.2 | Configure BGP router-id. |
| PE2-ASBR(config-router)# no bgp default ipv4-unicast | Configure BGP to disable default exchange of ipv4 unicast AFI/SAFI capability. |
| PE2-ASBR (config-router)# address-family ipv4 unicast | Enter address-family ipv4 unicast mode |
| PE2-ASBR(config-router-af)# redistribute ospf route-map RR1_LO | Redistribute RR1 Lo from OSPF to BGP. This is required for eBGP session between RR1 and RR2. |
| PE2-ASBR(config-router-af)#exit | Exit address-family mode |
| PE2-ASBR(config-router)# neighbor 20.1.1.40 remote-as 65001 | Configure PE3-ASBR as eBGP neighbor. |
| PE2-ASBR(config-router)# neighbor 20.1.1.40 activate | Activate PE3-ASBR neighbor for IPv4 unicast AFI/SAFI. This is required to exchange the redistributed RR1 Lo prefix to PE3-ASBR eBGP peer. |
| PE2-ASBR(config-router)# neighbor 31.31.31.31 remote-as 65000 | Configure RR1 as an iBGP neighbor. |
| PE2-ASBR(config-router)# neighbor 31.31.31.31 update-source lo | Update the source of the iBGP session to loopback address. |
| PE2-ASBR(config-router)# allocate-label all | Configure this command to allocate label for all advertised networks. |
| PE2-ASBR(config-router)# address-family ipv4 labeled-unicast | Enter ipv4 labeled-unicast AFI/SAFI. |
| PE2-ASBR(config-router-af)# neighbor 20.1.1.40 activate | Activate PE3-ASBR neighbor for IPv4 labeled-unicast. |
| PE2-ASBR(config-router-af)# neighbor 31.31.31.31 activate | Activate RR1 neighbor for IPv4 labeled-unicast. |
| PE2-ASBR(config-router-af)# neighbor 31.31.31.31 next-hop-self | Configure next-hop-self for the RR1 neighbor. |
| PE2-ASBR(config-router-af)# exit-address-family | Exit this address-family. |
| PE2-ASBR(config-router)#exit | Exit the router mode. |
| PE2-ASBR (config)#commit | Commit candidate configuration to be running |

| | |
|-----------------------|-----------------------|
| | configuration. |
| PE2-ASBR(config)#exit | Exit the config mode. |

PE3-ASBR

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| PE3-ASBR(config)#interface lo | Enter interface mode. |
| PE3-ASBR(config-if)# ip address 40.40.40.40/32 secondary | Assign a secondary IPv4 address to the loopback interface. |
| PE3-ASBR(config-if)#interface xe5 | Enter interface mode. |
| PE3-ASBR(config-if)# ip address 20.1.1.40/24 | Assign an IPv4 address to the interface. |
| PE3-ASBR(config-if)#interface xe1 | Enter interface mode. |
| PE3-ASBR(config-if)# ip address 21.1.1.40/24 | Assign an IPv4 address to the interface. |
| PE3-ASBR(config-if)#ip access-list standard RR1_LO | Create an IP standard access list to permit RR1s loopback address. This is required to redistribute RR1 Lo from BGP to OSPF. |
| PE3-ASBR(config-ip-acl-std)# permit host 31.31.31.31 | Permit RR1 loopback address. |
| PE3-ASBR(config-ip-acl-std)#ip access-list standard RR2_LO | Create an IP standard access list to permit RR2 Lo address. This is required to redistribute RR2 Lo from OSPF to BGP. |
| PE3-ASBR(config-ip-acl-std)# permit host 5.5.5.5 | Permit RR2s Loopback address. |
| PE3-ASBR(config-ip-acl-std)#route-map RR1_LO permit 10 | Create a Route-map to allow RR1-LO address . This is required when redistributing from BGP to OSPF. |
| PE3-ASBR(config-route-map)# match ip address RR1_LO | match the RR1_LO access list. |
| PE3-ASBR(config-route-map)#route-map RR2_LO permit 10 | Create a Route-map to allow RR2-LO address . This is required when redistributing from OSPF to BGP. |
| PE3-ASBR(config-route-map)# match ip address RR2_LO | match the RR2_LO access list. |
| PE3-ASBR(config-route-map)#router ospf 1 | Enter router OSPF mode. |
| PE3-ASBR(config-router)# ospf router-id 40.40.40.40 | Configure OSPF router-id explicitly (optional). |
| PE3-ASBR(config-router)# redistribute bgp route-map RR1_LO | Redistribute RR2's Lo from OSPF to BGP. This is required for eBGP session between RR1 and RR2. |
| PE3-ASBR(config-router)# network 21.1.1.0/24 area 0.0.0.0 | Enable OSPF networks/links in area 0. |
| PE3-ASBR(config-router)# network 40.40.40.40/32 area 0.0.0.0 | Enable OSPF networks/links in area 0. |
| PE3-ASBR(config-router)#router ldp | Enter Router LDP mode. This is required to enable LDP globally on the router. |
| PE3-ASBR(config-router)#router-id 40.40.40.40 | Set the router ID to IP address 40.40.40.40 |
| PE3-ASBR(config-router)#transport-address ipv4 40.40.40.40 | Configure LDP transport address same as loopback address |

| | |
|--|---|
| PE3-ASBR(config-router)#interface xe5 | Enter interface mode. |
| PE3-ASBR(config-if)# label-switching | Configure label-switching to enable MPLS forwarding on this interface. |
| PE3-ASBR(config-if)#interface xe1 | Enter interface mode. |
| PE3-ASBR(config-if)# label-switching | Configure label-switching to enable MPLS forwarding on this interface. |
| PE3-ASBR(config-if)# enable-ldp ipv4 | Enable LDP on this interface. |
| PE3-ASBR(config-if)#router bgp 65001 | Enter router BGP configuration mode. |
| PE3-ASBR(config-router)# bgp router-id 40.40.40.40 | Configure BGP router-id. |
| PE3-ASBR(config-router)# no bgp default ipv4-unicast | Configure BGP to disable default exchange of ipv4 unicast AFI/SAFI capability. |
| PE3-ASBR(config-router)# redistribute ospf route-map RR2_LO | Redistribute RR2 Lo from OSPF to BGP. This is required for eBGP session between RR1 and RR2. |
| PE3-ASBR(config-router)# neighbor 5.5.5.5 remote-as 65001 | Configure RR2 as eBGP neighbor. |
| PE3-ASBR(config-router)# neighbor 5.5.5.5 update-source lo | Update the source of the iBGP session to loopback address. |
| PE3-ASBR(config-router)# neighbor 20.1.1.2 remote-as 65000 | Configure PE2-ASBR as an eBGP neighbor. |
| PE3-ASBR(config-router)# neighbor 20.1.1.2 activate | Activate PE2-ASBR neighbor for IPv4 unicast AFI/SAFI. This is required to exchange the redistributed RR1 Lo prefix to PE2-ASBR eBGP peer. |
| PE3-ASBR(config-router)# allocate-label all | Configure this command to allocate label for all advertised networks. |
| PE3-ASBR(config-router)# address-family ipv4 labeled-unicast | Enter ipv4 labeled-unicast AFI/SAFI. |
| PE3-ASBR(config-router-af)# neighbor 5.5.5.5 activate | Activate RR1 neighbor for IPv4 labeled-unicast. |
| PE3-ASBR(config-router-af)# neighbor 5.5.5.5 next-hop-self | Configure next-hop-self for the RR2 neighbor. |
| PE3-ASBR(config-router-af)# neighbor 20.1.1.2 activate | Activate PE2-ASBR neighbor for IPv4 labeled-unicast. |
| PE3-ASBR(config-router-af)# ex-it-address-family | Exit this address-family. |
| PE3-ASBR(config-router)#exit | Exit the router mode. |
| PE3-ASBR(config)#commit | Commit candidate configuration to be running configuration. |
| PE3-ASBR(config)#exit | Exit the config mode. |

RR2

| | |
|--------------------------|-----------------------|
| #configure terminal | Enter configure mode. |
| RR2(config)#interface lo | Enter interface mode. |

| | |
|--|--|
| RR2(config-if)# ip address 5.5.5.5/32 secondary | Assign the IPv4 address. |
| RR2(config-if)#interface xe22 | Enter interface mode. |
| RR2(config-if)# ip address 23.1.1.2/24 | Assign the IPv4 address. |
| RR2(config-if)#router ospf 1 | Enter router OSPF mode. |
| RR2(config-router)# ospf router-id 5.5.5.5 | Configure OSPF router-id explicitly (optional). |
| RR2(config-router)# network 5.5.5.5/32 area 0.0.0.0 | Enable networks for OSPF protocol under area 0. |
| RR2(config-router)# network 23.1.1.0/24 area 0.0.0.0 | Enable networks for OSPF protocol under area 0. |
| RR2(config-router)#router bgp 65001 | Enter BGP router mode. |
| RR2(config-router)# bgp router-id 5.5.5.5 | Configure BGP router-id. |
| RR2(config-router)# no bgp default ipv4-unicast | Disable default ipv4-unicast capability exchange with BGP neighbors. |
| RR2(config-router)# no bgp inbound-route-filter | Disable inbound route-filtering for bgp VPNV4 routes. This is required for Route-reflectors. |
| RR2(config-router)# neighbor 28.28.28.28 remote-as 65001 | Configure PE2-ASBR as iBGP neighbor. |
| RR2(config-router)# neighbor 28.28.28.28 update-source lo | Configure BGP speaker to use loopback address as source for BGP session. |
| RR2(config-router)# neighbor 31.31.31.31 remote-as 65000 | Configure RR2 as an eBGP neighbor. |
| RR2(config-router)# neighbor 31.31.31.31 ebgp-multihop | Enable multihop on eBGP session. |
| RR2(config-router)# neighbor 31.31.31.31 update-source lo | Configure BGP speaker to use loopback address as source for this BGP session. |
| RR2(config-router)# neighbor 40.40.40.40 remote-as 65001 | Configure PE1 as an iBGP neighbor. |
| RR2(config-router)# neighbor 40.40.40.40 update-source lo | Configure BGP speaker to use loopback address as source for this BGP session. |
| RR2(config-router)# allocate-label all | Configure this peer to allocate label for all advertised networks. |
| RR2(config-router)# address-family ipv4 labeled-unicast | Enter ipv4 labeled-unicast address-family. |
| RR2(config-router-af)# neighbor 28.28.28.28 activate | Activate PE2-ASBR for this AFI/SAFI. |
| RR2(config-router-af)# neighbor 28.28.28.28 route-reflector-client | Configure PE2-ASBR as Route-reflector client. |
| RR2(config-router-af)# neighbor 40.40.40.40 activate | Activate PE1 for this AFI/SAFI. |
| RR2(config-router-af)# neighbor 40.40.40.40 route-reflector-client | Configure PE1 as Route-reflector client. |
| RR2(config-router-af)# exit-address-family | Exit ipv4 labeled-unicast address-family. |
| RR2(config-router)# address-family vpnv4 unicast | Enter vpnv4 unicast address-family. |
| RR2(config-router-af)# neighbor 28.28.28.28 | Allow eBGP vpn route exchange for the RR2 eBGP |

| | |
|--|--|
| activate | neighbor. |
| RR2(config-router-af)# neighbor 28.28.28.28 route-reflector-client | Activate the RR2 eBGP neighbor. |
| RR2(config-router-af)# neighbor 31.31.31.31 allow-ebgp-vpn | Configure next-hop unchanged for this eBGP neighbor for all vpnv4 NLRIs to keep original next-hop while advertising to this eBGP neighbor. |
| RR2(config-router-af)# neighbor 31.31.31.31 activate | Activate PE1 iBGP neighbor. |
| RR2(config-router-af)# neighbor 31.31.31.31 attribute-unchanged next-hop | Configure PE1 as Route-reflector client for this AFI/SAFI. |
| RR2(config-router-af)# exit-address-family | Exit Address-family mode. |
| RR2(config-router)#exit | Exit the router mode. |
| RR2(config)#commit | Commit candidate configuration to be running configuration. |
| RR2(config)#exit | Exit the config mode. |

P2

| | |
|---|---|
| P2(config)#interface lo | Enter interface mode. |
| P2(config-if)# ip address 6.6.6.6/32 secondary | Assign the secondary IPv4 address on loopback interface. |
| P2(config-if)#interface xe22 | Enter interface mode. |
| P2(config-if)# ip address 23.1.1.1/24 | Assign the IPv4 address. |
| P2(config-if)#interface xe0 | Enter interface mode. |
| P2(config-if)# ip address 21.1.1.5/24 | Assign the IPv4 address. |
| P2(config-if)#interface xe11 | Enter interface mode. |
| P2(config-if)# ip address 22.1.1.5/24 | Assign the IPv4 address. |
| P2(config-if)#router ospf 1 | Enter router OSPF mode. |
| P2(config-router)# ospf router-id 6.6.6.6 | Configure OSPF router-id (optional). |
| P2(config-router)# network 6.6.6.6/32 area 0.0.0.0 | Enable networks for OSPF protocol. |
| P2(config-router)# network 21.1.1.0/24 area 0.0.0.0 | Enable networks for OSPF protocol. |
| P2(config-router)# network 22.1.1.0/24 area 0.0.0.0 | Enable networks for OSPF protocol. |
| P2(config-router)# network 23.1.1.0/24 area 0.0.0.0 | Enable networks for OSPF protocol. |
| P2(config-router)#router ldp | Enter router ldp mode. This is required to enable LDP globally on the router. |
| P2(config-router)#router-id 6.6.6.6 | Set the router ID to IP address 6.6.6.6 |
| P2(config-router)#transport-address ipv4 6.6.6.6 | Configure LDP transport address same as loopback address. |
| P2(config-router)#interface xe0 | Enter interface configuration mode. |
| P2(config-if)# label-switching | Configure label-switching on the interface. |

| | |
|--------------------------------|---|
| P2(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| P2(config-if)#interface xe11 | Enter interface configuration mode. |
| P2(config-if)# label-switching | Configure label-switching on the interface. |
| P2(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| P2(config-if)#exit | Exit interface mode. |
| P2(config)#commit | Commit candidate configuration to be running configuration. |
| P2(config)#exit | Exit the config mode. |

PE4

| | |
|---|--|
| #configure terminal | Enter configure mode. |
| PE4(config)#interface lo | Enter interface mode. |
| PE4(config-if)# ip address 28.28.28.28/32 secondary | Assign the secondary IPv4 address on the loopback interface. |
| PE4(config-if)#ip vrf vrf1 | Create a new VRF named vrf1. |
| PE4(config-vrf)# rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| PE4(config-vrf)# route-target both 1:1 | Configure import and export route-target values. |
| PE4(config-vrf)#interface xe21 | Enter interface mode. |
| PE4(config-if)# ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1. |
| PE4(config-if)# ip address 30.1.1.1/24 | Assign the IPv4 address. |
| PE4(config-if)#interface xe11 | Enter interface mode. |
| PE4(config-if)# ip address 22.1.1.28/24 | Assign the IPv4 address. |
| PE4(config-if)#router ospf 1 | Enter router OSPF mode. |
| PE4(config-router)# ospf router-id 28.28.28.28 | Configure OSPF router id same as loopback ip address. |
| PE4(config-router)# network 22.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE4(config-router)# network 28.28.28.28/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE4(config-router)#router ldp | Enter router LDP mode. |
| PE4(config-router)#router-id 28.28.28.28 | Set the router ID to IP address 28.28.28.28 |
| PE4(config-router)#transport-address ipv4 28.28.28.28 | Configure LDP transport address same as loopback address. |
| PE4(config-router)#interface xe11 | Enter interface mode. |
| PE4(config-if)# label-switching | Enable label switching in interface. |
| PE4(config-if)# enable-ldp ipv4 | Enable LDP in interface. |
| PE4(config-if)#router bgp 65001 | Enter BGP router mode. |

| | |
|---|--|
| PE4(config-router)# bgp router-id 28.28.28.28 | Configure BGP router-id. |
| PE4(config-router)# no bgp default ipv4-unicast | Configure BGP peer to have no ipv4 unicast capability exchange by default. |
| PE4(config-router)# address-family ipv4 unicast | Enter the IPv4 unicast address family. |
| PE4(config-router-af)# network 28.28.28.28/32 | advertise the network of this BGP speaker. |
| PE4(config-router-af)# exit-address-family | Exit address family mode. |
| PE4(config-router)# neighbor 5.5.5.5 remote-as 65001 | configure neighbor with remote AS. |
| PE4(config-router)# neighbor 5.5.5.5 update-source lo | Update the loopback as the source of BGP session. |
| PE4(config-router)# allocate-label all | Allocate the label for all advertised networks. |
| PE4(config-router)# address-family ipv4 labeled-unicast | Enter the IPv4 labeled-unicast address family. |
| PE4(config-router-af)# neighbor 5.5.5.5 activate | Activate the neighbor under the address-family. |
| PE4(config-router-af)# exit-address-family | Exit IPv4 labeled-unicast Address Family mode. |
| PE4(config-router)# address-family vpnv4 unicast | Enter vpnv4 unicast address-family mode. |
| PE4(config-router-af)# neighbor 5.5.5.5 activate | Activate the neighbor under vpnv4 unicast address-family. |
| PE4(config-router-af)# exit-address-family | Exit vpnv4 unicast address-family. |
| PE4(config-router)# address-family ipv4 vrf vrf1 | Enter ipv4 VRF address-family mode. |
| PE4(config-router-af)# redistribute connected | Redistribute connected networks under VRF address-family. |
| PE4(config-router-af)# exit-address-family | Exit ipv4 VRF address-family. |
| PE4(config-router)#exit | Exit the router mode. |
| PE4(config)#commit | Commit candidate configuration to be running configuration. |
| PE4(config)#exit | Exit the config mode. |

Validation

PE1

```

PE1#show ip bgp labeled-unicast summary
BGP router identifier 41.41.41.41, local AS number 65000
BGP table version is 5
2 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS    MsgRcv   MsgSen  TblVer   InQ    OutQ    Up/Down   State/PfxRcd
31.31.31.31        4  65000  5920     5932      5        0        0  1d15h58m        3

Total number of neighbors 1

Total number of Established sessions 1

PE1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
Network          Next Hop        In Label      Out Label

```

```
*>i 5.5.5.5/32      2.2.2.2      24963      24965
*>i 28.28.28.28/32  2.2.2.2      24965      24964
*>i 31.31.31.31/32  2.2.2.2      24967      24961
*> 41.41.41.41/32  0.0.0.0      24961      -
```

PE1#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|----------------|---------|-----------|-----------|-------------|-----------|------|
| Intf | ELC | Nexthop | | | | | |
| L> | 2.2.2.2/32 | 4 | 0 | Yes | LSP_ | | |
| DEFAULT | 24321 | xe15 | No | 11.1.1.31 | | | |
| L> | 5.5.5.5/32 | 6 | 0 | Yes | LSP_ | | |
| DEFAULT | 24324 | xe15 | No | 11.1.1.31 | | | |
| B | 5.5.5.5/32 | 7 | 0 | Yes | LSP_DEFAULT | 24965 | - |
| | No | 2.2.2.2 | | | | | |
| L> | 12.1.1.0/24 | 2 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe15 | No | 11.1.1.31 | | | |
| L> | 13.1.1.0/24 | 3 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | e15 | No | 11.1.1.31 | | | |
| B> | 28.28.28.28/32 | 9 | 0 | Yes | LSP_DEFAULT | 24964 | - |
| | No | 2.2.2.2 | | | | | |
| L> | 31.31.31.31/32 | 5 | 0 | Yes | LSP_ | | |
| DEFAULT | 24322 | xe15 | No | 11.1.1.31 | | | |
| B | 31.31.31.31/32 | 11 | 0 | Yes | LSP_DEFAULT | 24961 | - |
| | No | 2.2.2.2 | | | | | |
| L> | 51.51.51.51/32 | 1 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe15 | No | 11.1.1.31 | | | |

PE1#show ip bgp vpnv4 all summary

BGP router identifier 41.41.41.41, local AS number 65000

BGP table version is 2

2 BGP AS-PATH entries

0 BGP community entries

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|-------------|---|-------|--------|--------|--------|-----|------|----------|--------------|
| 31.31.31.31 | 4 | 65000 | 5928 | 5940 | 2 | 0 | 0 | 1d16h01m | 1 |

Total number of neighbors 1

Total number of Established sessions 1

PE1#show ip bgp vpnv4 all

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
 S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|---|-------------|--------|--------|--------|---------|
| Route Distinguisher: 1:1 (Default for VRF vrf1) | | | | | |
| *> 1 10.1.1.0/24 | 0.0.0.0 | 0 | 100 | 32768 | ? |
| *>i 1 30.1.1.0/24 | 28.28.28.28 | 0 | 100 | 0 | 65001 ? |
| Announced routes count = 1 | | | | | |
| Accepted routes count = 1 | | | | | |
| Route Distinguisher: 1:1 | | | | | |
| *>i 1 30.1.1.0/24 | 28.28.28.28 | 0 | 100 | 0 | 65001 ? |
| Announced routes count = 0 | | | | | |
| Accepted routes count = 1 | | | | | |

PE1#show ip route vrf vrf1 database

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP

O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,

ia - IS-IS inter area, E - EVPN,

v - vrf leaked


```

    > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1"
C    *> 10.1.1.0/24 is directly connected, xe23, 1d15h59m
B    *> 30.1.1.0/24 [200/0] via 28.28.28.28, 1d15h47m

Gateway of last resort is not set

```

P1

```

P1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, S - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

Code   FEC/VRF/L2CKT   ILM-ID   In-Label   Out-Label   In-Intf   Out-
Intf/VRF      Nexthop      LSP-Type
L>    2.2.2.2/32      2         24321      3           N/A       xe21      12.1.1.2
      LSP_DEFAULT
L>    41.41.41.41/32  1         24320      3           N/A       xe15      11.1.1.2
      LSP_DEFAULT
L>    31.31.31.31/32  4         24323      Nolabel     N/A       N/A       127.0.0.1
      LSP_DEFAULT
L>    31.31.31.31/32  3         24322      Nolabel     N/A       N/A       127.0.0.1
      LSP_DEFAULT
L>    5.5.5.5/32      5         24324      Nolabel     N/A       N/A       127.0.0.1
      LSP_DEFAULT

```

RR1

```

RR1#show ip route ospf
IP Route Table for VRF "default"
O      2.2.2.2/32 [110/3] via 13.1.1.1, xe22, 1d16h37m
O E2    5.5.5.5/32 [110/1] via 13.1.1.1, xe22, 1d16h35m
O      11.1.1.0/24 [110/2] via 13.1.1.1, xe22, 1d16h37m
O      12.1.1.0/24 [110/2] via 13.1.1.1, xe22, 1d16h37m
O      41.41.41.41/32 [110/3] via 13.1.1.1, xe22, 1d16h37m
O      51.51.51.51/32 [110/2] via 13.1.1.1, xe22, 1d16h37m

Gateway of last resort is not set

RR1#sho ip bgp labeled-unicast summary
BGP router identifier 31.31.31.31, local AS number 65000
BGP table version is 9
2 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv  MsgSen  TblVer  InQ   OutQ   Up/Down  State/PfxRcd
2.2.2.2            4  65000  7514    7654     9      0      0  1d16h11m      3
41.41.41.41        4  65000  5961    5951     9      0      0  1d16h11m      1

Total number of neighbors 2

Total number of Established sessions 2

RR1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
   Network        Next Hop        In Label        Out Label
*>i 5.5.5.5/32     2.2.2.2         24320           24965
*>i 28.28.28.28/32 2.2.2.2         24323           24964
*>i 31.31.31.31/32 2.2.2.2         24327           24961
*>i 41.41.41.41/32 41.41.41.41     24326           24961

```

```
RR1#show ip bgp vpnv4 all summary
BGP router identifier 31.31.31.31, local AS number 65000
BGP table version is 3
2 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|-------------|---|-------|--------|--------|--------|-----|------|----------|--------------|
| 5.5.5.5 | 4 | 65001 | 5729 | 5725 | 3 | 0 | 0 | 1d15h56m | 1 |
| 41.41.41.41 | 4 | 65000 | 5962 | 5953 | 3 | 0 | 0 | 1d16h12m | 1 |

```
Total number of neighbors 2

Total number of Established sessions 2

RR1#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|--------------------------|-------------|--------|--------|--------|---------|
| Route Distinguisher: 1:1 | | | | | |
| *>i 10.1.1.0/24 | 41.41.41.41 | 0 | 100 | 0 | ? |
| *> 30.1.1.0/24 | 28.28.28.28 | 0 | 100 | 0 | 65001 ? |

```
Announced routes count = 0
Accepted routes count = 2
```

PE2-ASBR

```
PE2-ASBR#show ip bgp labeled-unicast summary
BGP router identifier 2.2.2.2, local AS number 65000
BGP table version is 6
2 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|-------------|---|-------|--------|--------|--------|-----|------|----------|--------------|
| 20.1.1.40 | 4 | 65001 | 5884 | 5976 | 6 | 0 | 0 | 1d16h37m | 2 |
| 31.31.31.31 | 4 | 65000 | 5790 | 5794 | 6 | 0 | 0 | 1d16h15m | 1 |

```
Total number of neighbors 2

Total number of Established sessions 2

PE2-ASBR#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
```

| Network | Next Hop | In Label | Out Label |
|--------------------|-------------|----------|-----------|
| *> 5.5.5.5/32 | 20.1.1.40 | 24965 | 24965 |
| *> 28.28.28.28/32 | 20.1.1.40 | 24964 | 24962 |
| *> 31.31.31.31/32 | 12.1.1.31 | 24961 | - |
| *>i 41.41.41.41/32 | 41.41.41.41 | 24967 | 24961 |

```
PE2-ASBR#sho mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, S - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS_SR, k - SR CLI
       P - SR Policy, U - unknown
```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|-------------------|---------------|--------|----------|-----------|---------|-----------|
| Intf/VRF | Nexthop | | LSP-Type | | | |
| B> 5.5.5.5/32 | 6 | 24965 | 24965 | N/A | N/A | 20.1.1.40 |
| | LSP_DEFAULT | | | | | |
| B> 28.28.28.28/32 | 5 | 24964 | 24962 | N/A | N/A | 20.1.1.40 |
| | LSP_DEFAULT | | | | | |
| B> 31.31.31.31/32 | 2 | 24961 | Nolabel | N/A | N/A | 127.0.0.1 |
| | LSP_DEFAULT | | | | | |

```

B> 41.41.41.41/32      8          24967      24961      N/A      N/A      41.41.41.4
1          LSP_DEFAULT

PE2-ASBR#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code   FEC          FTN-ID   Tunnel-id   Pri   LSP-Type   Out-Label   Out-
Intf   ELC      Nextthop
B> 5.5.5.5/32      8          0          Yes   LSP_
DEFAULT 24965      xe5        No      20.1.1.40
L> 11.1.1.0/24     1          0          Yes   LSP_
DEFAULT 3          xe20       No      12.1.1.31
L> 13.1.1.0/24     2          0          Yes   LSP_
DEFAULT 3          xe20       No      12.1.1.31
B> 28.28.28.28/32  7          0          Yes   LSP_
DEFAULT 24962      xe5        No      20.1.1.40
L> 31.31.31.31/32  5          0          Yes   LSP_
DEFAULT 24323      xe20       No      12.1.1.31
L> 41.41.41.41/32  3          0          Yes   LSP_
DEFAULT 24320      xe20       No      12.1.1.31
B 41.41.41.41/32  10         0          Yes   LSP_DEFAULT 24961      -
      No      41.41.41.41
L> 51.51.51.51/32  4          0          Yes   LSP_
DEFAULT 3          xe20       No      12.1.1.31

PE2-ASBR#show ip route bgp
IP Route Table for VRF "default"
B      5.5.5.5/32 [20/3] via 20.1.1.40, xe5, 1d16h38m
B      28.28.28.28/32 [20/0] via 20.1.1.40, xe5, 1d16h39m

Gateway of last resort is not set

PE2-ASBR#show ip route ospf
IP Route Table for VRF "default"
O      11.1.1.0/24 [110/2] via 12.1.1.31, xe20, 1d16h59m
O      13.1.1.0/24 [110/2] via 12.1.1.31, xe20, 1d16h59m
O      31.31.31.31/32 [110/3] via 12.1.1.31, xe20, 1d16h41m
O      41.41.41.41/32 [110/3] via 12.1.1.31, xe20, 1d16h59m
O      51.51.51.51/32 [110/2] via 12.1.1.31, xe20, 1d16h59m

Gateway of last resort is not set

PE2-ASBR#show ip route connected
IP Route Table for VRF "default"
C      2.2.2.2/32 is directly connected, lo, 1d18h27m
C      12.1.1.0/24 is directly connected, xe20, 1d18h27m
C      20.1.1.0/24 is directly connected, xe5, 1d18h27m
C      127.0.0.0/8 is directly connected, lo, 1d19h32m

Gateway of last resort is not set

```

PE3-ASBR

```

PE3-ASBR#show ip bgp labeled-unicast summary
BGP router identifier 40.40.40.40, local AS number 65001
BGP table version is 7
2 BGP AS-PATH entries
0 BGP community entries

Neighbor      V  AS  MsgRcv  MsgSen  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
5.5.5.5       4 65001 5729    5736    7       0     0  1d16h46m 1
20.1.1.2      4 65000 5731    5739    7       0     0  1d16h46m 2

Total number of neighbors 2

```

Total number of Established sessions 2

PE3-ASBR#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale

| Network | Next Hop | In Label | Out Label |
|--------------------|-------------|----------|-----------|
| *> 5.5.5.5/32 | 21.1.1.5 | 24965 | - |
| *>i 28.28.28.28/32 | 28.28.28.28 | 24962 | 24321 |
| *> 31.31.31.31/32 | 20.1.1.2 | 24964 | 24961 |
| *> 41.41.41.41/32 | 20.1.1.2 | 24967 | 24967 |

PE3-ASBR#sho mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM
 K - CLI ILM, T - MPLS-TP, S - Stitched ILM
 S - SNMP, L - LDP, R - RSVP, C - CRLDP
 B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
 O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
 P - SR Policy, U - unknown

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|----------------|--------|----------|-----------|---------|------------|
| Intf/VRF | NextHop | | LSP-Type | | | |
| B> | 31.31.31.31/32 | 6 | 24964 | 24961 | N/A | N/A |
| | LSP_DEFAULT | | | | | 20.1.1.2 |
| B> | 28.28.28.28/32 | 4 | 24962 | 24321 | N/A | N/A |
| 8 | LSP_DEFAULT | | | | | 28.28.28.2 |
| B> | 41.41.41.41/32 | 9 | 24967 | 24967 | N/A | N/A |
| | LSP_DEFAULT | | | | | 20.1.1.2 |
| B> | 5.5.5.5/32 | 7 | 24965 | Nolabel | N/A | N/A |
| | LSP_DEFAULT | | | | | 127.0.0.1 |

PE3-ASBR#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|----------------|-------------|-----------|----------|-------------|-----------|------|
| Intf | ELC | NextHop | | | | | |
| L> | 5.5.5.5/32 | 5 | 0 | Yes | LSP_ | | |
| DEFAULT | 24320 | xe1 | No | 21.1.1.5 | | | |
| L> | 6.6.6.6/32 | 6 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe1 | No | 21.1.1.5 | | | |
| L> | 22.1.1.0/24 | 7 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe1 | No | 21.1.1.5 | | | |
| L> | 23.1.1.0/24 | 8 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe1 | No | 21.1.1.5 | | | |
| L> | 28.28.28.28/32 | 9 | 0 | Yes | LSP_ | | |
| DEFAULT | 24321 | xe1 | No | 21.1.1.5 | | | |
| B | 28.28.28.28/32 | 2 | 0 | Yes | LSP_DEFAULT | 24321 | - |
| | No | 28.28.28.28 | | | | | |
| B> | 31.31.31.31/32 | 4 | 0 | Yes | LSP_ | | |
| DEFAULT | 24961 | xe5 | No | 20.1.1.2 | | | |
| B> | 41.41.41.41/32 | 11 | 0 | Yes | LSP_ | | |
| DEFAULT | 24967 | xe5 | No | 20.1.1.2 | | | |

PE3-ASBR#show ip route bgp

IP Route Table for VRF "default"

B 31.31.31.31/32 [20/3] via 20.1.1.2, xe5, 1d16h46m
 B 41.41.41.41/32 [20/0] via 20.1.1.2, xe5, 1d16h23m

Gateway of last resort is not set

PE3-ASBR#show ip route ospf

IP Route Table for VRF "default"

O 5.5.5.5/32 [110/3] via 21.1.1.5, xe1, 1d16h54m
 O 6.6.6.6/32 [110/2] via 21.1.1.5, xe1, 1d16h55m
 O 22.1.1.0/24 [110/2] via 21.1.1.5, xe1, 1d16h55m
 O 23.1.1.0/24 [110/2] via 21.1.1.5, xe1, 1d16h54m
 O 28.28.28.28/32 [110/3] via 21.1.1.5, xe1, 1d16h55m

```

PE3-ASBR#show ip route connected
IP Route Table for VRF "default"
C          20.1.1.0/24 is directly connected, xe5, 1d16h55m
C          21.1.1.0/24 is directly connected, xe1, 1d16h55m
C          40.40.40.40/32 is directly connected, lo, 1d16h55m
C          127.0.0.0/8 is directly connected, lo, 1d19h39m

Gateway of last resort is not set

```

RR2

```

RR2#show ip bgp labeled-unicast summary
BGP router identifier 5.5.5.5, local AS number 65001
BGP table version is 10
2 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS    MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd
28.28.28.28        4  65001  5795     5815    10        0     0    1d17h00m        1
40.40.40.40        4  65001  5779     5787    10        0     0    1d16h54m        3

Total number of neighbors 2

Total number of Established sessions 2

RR2#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
   Network          Next Hop          In Label    Out Label
*>i 5.5.5.5/32       40.40.40.40       24325       24965
*>i 28.28.28.28/32   28.28.28.28       24322       24321
*>i 31.31.31.31/32   40.40.40.40       24324       24964
*>i 41.41.41.41/32   40.40.40.40       24327       24967

RR2#show ip bgp vpnv4 all summary
BGP router identifier 5.5.5.5, local AS number 65001
BGP table version is 3
2 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS    MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd
28.28.28.28        4  65001  5796     5815     3        0     0    1d17h00m        1
31.31.31.31        4  65000  5769     5776     3        0     0    1d16h16m        1

Total number of neighbors 2

Total number of Established sessions 2

RR2#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric    LocPrf    Weight Path
Route Distinguisher: 1:1
*>  10.1.1.0/24       41.41.41.41         0         100         0       65000 ?
*>i 30.1.1.0/24       28.28.28.28         0         100         0         ?
Announced routes count = 0
Accepted routes count = 2

RR2#show ip route ospf
IP Route Table for VRF "default"
O          6.6.6.6/32 [110/2] via 23.1.1.1, xe22, 1d17h02m
O          21.1.1.0/24 [110/2] via 23.1.1.1, xe22, 1d17h02m
O          22.1.1.0/24 [110/2] via 23.1.1.1, xe22, 1d17h02m
O          28.28.28.28/32 [110/3] via 23.1.1.1, xe22, 1d17h02m
O E2       31.31.31.31/32 [110/1] via 23.1.1.1, xe22, 1d16h54m

```

```
O          40.40.40.40/32 [110/3] via 23.1.1.1, xe22, 1d17h02m
```

```
Gateway of last resort is not set
```

P2

```
P2#show mpls ilm-table
```

```
Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM, T - MPLS-TP, S - Stitched ILM
        S - SNMP, L - LDP, R - RSVP, C - CRLDP
        B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
        O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
        P - SR Policy, U - unknown
```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|----------------|--------|----------|-----------|---------|-----------|
| Intf/VRF | Nextthop | | LSP-Type | | | |
| L> | 5.5.5.5/32 | 3 | 24322 | Nolabel | N/A | N/A |
| | LSP_DEFAULT | | | | | 127.0.0.1 |
| L> | 5.5.5.5/32 | 1 | 24320 | Nolabel | N/A | N/A |
| | LSP_DEFAULT | | | | | 127.0.0.1 |
| L> | 28.28.28.28/32 | 2 | 24321 | 3 | N/A | xe11 |
| | LSP_DEFAULT | | | | | 22.1.1.28 |
| L> | 31.31.31.31/32 | 4 | 24323 | Nolabel | N/A | N/A |
| | LSP_DEFAULT | | | | | 127.0.0.1 |
| L> | 40.40.40.40/32 | 5 | 24324 | 3 | N/A | xe0 |
| | LSP_DEFAULT | | | | | 21.1.1.40 |

PE4

```
PE4#show ip bgp labeled-unicast
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
```

| Network | Next Hop | In Label | Out Label |
|--------------------|-------------|----------|-----------|
| *>i 5.5.5.5/32 | 40.40.40.40 | 24325 | 24965 |
| *> 28.28.28.28/32 | 0.0.0.0 | 24321 | - |
| *>i 31.31.31.31/32 | 40.40.40.40 | 24324 | 24964 |
| *>i 41.41.41.41/32 | 40.40.40.40 | 24327 | 24967 |

```
PE4#show mpls forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
        B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
```

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|----------------|-------------|-----------|----------|-------------|-----------|------|
| Intf | ELC | Nextthop | | | | | |
| L> | 5.5.5.5/32 | 4 | 0 | Yes | LSP_ | | |
| DEFAULT | 24322 | xe11 | No | 22.1.1.5 | | | |
| B | 5.5.5.5/32 | 3 | 0 | Yes | LSP_DEFAULT | 24965 | - |
| | No | 40.40.40.40 | | | | | |
| L> | 6.6.6.6/32 | 5 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe11 | No | 22.1.1.5 | | | |
| L> | 21.1.1.0/24 | 6 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe11 | No | 22.1.1.5 | | | |
| L> | 23.1.1.0/24 | 7 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe11 | No | 22.1.1.5 | | | |
| L> | 31.31.31.31/32 | 8 | 0 | Yes | LSP_ | | |
| DEFAULT | 24323 | xe11 | No | 22.1.1.5 | | | |
| B | 31.31.31.31/32 | 2 | 0 | Yes | LSP_DEFAULT | 24964 | - |
| | No | 40.40.40.40 | | | | | |
| L> | 40.40.40.40/32 | 9 | 0 | Yes | LSP_ | | |
| DEFAULT | 24324 | xe11 | No | 22.1.1.5 | | | |
| B> | 41.41.41.41/32 | 11 | 0 | Yes | LSP_DEFAULT | 24967 | - |
| | No | 40.40.40.40 | | | | | |

```
PE4#show ip bgp vpnv4 all summary
```

```
BGP router identifier 28.28.28.28, local AS number 65001
```

```
BGP table version is 2
```

```

2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V    AS    MsgRcv    MsgSen    TblVer    InQ    OutQ    Up/Down    State/PfxRcd
5.5.5.5                  4 65001 5792      5781      2        0        0    1d17h08m          1

Total number of neighbors 1

Total number of Established sessions 1

PE4#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric    LocPrf    Weight Path
Route Distinguisher: 1:1 (Default for VRF vrf1)
*>i 10.1.1.0/24      41.41.41.41         0         100        0        65000 ?
*> 1 30.1.1.0/24     0.0.0.0             0         100       32768      ?
  Announced routes count = 1
  Accepted routes count = 1
Route Distinguisher: 1:1
*>i 10.1.1.0/24      41.41.41.41         0         100        0        65000 ?
  Announced routes count = 0
  Accepted routes count = 1

PE4#show ip route vrf vrf1 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1"
B    *> 10.1.1.0/24 [200/0] via 41.41.41.41, 1d16h23m
C    *> 30.1.1.0/24 is directly connected, xe21, 1d16h35m

Gateway of last resort is not set

```

Configuration Without RR


This sub-option enables VPNv4 routes exchange between the Provider Edge Routers of different AS'es and thus requires

1. Remote PE loopback addresses to be known in the AS - This is accomplished using BGP-LU and is advertised inside the AS using the Route-reflector and outside the AS using the eBGP Peering between the ASBRs.
2. Next-hop attribute should be changed when ASBR advertises the remote PE's Loopback address as LU route to the RR. This is accomplished using the configuration "next-hop self" for the iBGP RR neighbor.
3. MP-eBGP session between the two PE's in different AS's to exchange the VPNv4 routes.

First and second points enables an end to end LSP between the two PE's which needs to run the L3VPN service and (3) enables the PE's in different AS'es to form BGP neighborhood with each other and successfully exchange the VPNv4 Routes.

Topology

Figure 92. InterAS-VPN Option-C without RR



PE1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| PE1(config)#interface lo | Enter interface mode. |
| PE1(config-if)#ip address 41.41.41.41/32 secondary | Assign the IPv4 address. |
| PE1(config-if)#exit | Exit interface mode. |
| PE1(config)#ip vrf vrf1 | Create a new VRF named vrf1. |
| PE1(config-vrf)#rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| PE1(config-vrf)#route-target both 1:1 | Configure import and export RT for this VRF. |
| PE1(config-vrf)#exit | Exit interface mode. |
| PE1(config)#interface xe23 | Enter interface mode. |
| PE1(config-if)#ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1. |
| PE1(config-if)#ip address 10.1.1.1/24 | Assign the IPv4 address. |
| PE1(config)#interface xe15 | Enter interface mode. |
| PE1(config-if)#ip address 11.1.1.41/24 | Assign the IPv4 address. |
| PE1(config-if)#exit | Exit interface mode. |
| PE1(config)#router ospf 1 | Enter router OSPF mode. |
| PE1(config-router)#ospf router-id 41.41.41.41 | Configure OSPF router id same as loopback ip address. |
| PE1(config-router)#network 11.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE1(config-router)#network 41.41.41.41/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE1(config-router)#exit | Exit OSPF router mode. |
| PE1(config-router)#router ldp | Enter router LDP mode. |
| PE1(config-router)#router-id 41.41.41.41 | Set the router ID to IP address 41.41.41.41 |
| PE1(config-router)#transport-address ipv4 41.41.41.41 | Configure LDP transport address same as loopback address |

| | |
|--|--|
| PE1(config-router)#exit | Exit LDP mode. |
| (config)#interface xe15 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#exit | Exit interface mode. |
| PE1(config-if)#router bgp 65000 | Enter BGP router mode. |
| PE1(config-router)# bgp router-id 41.41.41.41 | Configure BGP router-id. |
| PE1(config-router)# no bgp default ipv4-unicast | Configure BGP peer to have no ipv4 unicast capability exchange by default. |
| PE1(config-router)# address-family ipv4 unicast | Enter address-family ipv4 unicast. |
| PE1(config-router-af)# network 41.41.41.41/32 | Advertise loopback address. |
| PE1(config-router-af)# exit-address-family | Exit address-family ipv4 labeled-unicast. |
| PE1(config-router)# neighbor 2.2.2.2 re-mote-as 65000 | Configure PE2-ASBR1 as an iBGP peer. |
| PE1(config-router)# neighbor 2.2.2.2 up-date-source lo | Update the source as loopback for iBGP peering with the remote PE2 router. |
| PE1(config-router)# neighbor 28.28.28.28 remote-as 65001 | Configure PE4 as an eBGP peer. |
| PE1(config-router)# neighbor 28.28.28.28 ebgp-multihop | Configure eBGP multichip for eBGP peer PE4. |
| PE1(config-router)# neighbor 28.28.28.28 update-source lo | Update the source as loopback for eBGP peering with the remote PE4 router. |
| PE1(config-router)# allocate-label all | Configure allocate-label. |
| PE1(config-router)# address-family ipv4 la-beled-unicast | Enter address-family ipv4 labeled-unicast. |
| PE1(config-router-af)# neighbor 2.2.2.2 ac-tivate | Activate iBGP neighbor. |
| PE1(config-router-af)# exit-address-family | Exit address-family ipv4 labeled-unicast. |
| PE1(config-router)# address-family vpnv4 unicast | Enter address-family vpnv4. |
| PE1(config-router-af)# neighbor 28.28.28.28 allow-ebgp-vpn | Configure allow-ebgp-vpn for ebgp neighbor PE4. |
| PE1(config-router-af)# neighbor 28.28.28.28 activate | Activate eBGP neighbor PE4. |
| PE1(config-router-af)# exit-address-family | Exit address-family vpnv4. |
| PE1(config-router)# address-family ipv4 vrf vrf1 | Enter the IPv4 address family for VRF vrf1. |
| PE1(config-router-af)# redistribute con-nected | Redistribute connected route. |
| PE1(config-router-af)# exit-address-family | Exit IPv4 VRF Address Family mode. |
| PE1(config-router)#exit | Exit the router mode. |
| PE1(config)#commit | Commit candidate configuration to be running configuration. |
| PE1(config)#exit | Exit the config mode. |

P1

| | |
|---|--|
| P1#configure terminal | Enter configure mode. |
| P1(config)#interface lo | Enter interface mode. |
| P1(config-if)# ip address 51.51.51.51/32 secondary | Assign the IPv4 address. |
| P1(config-if)#exit | Exit interface mode. |
| P1(config)#interface xe15 | Enter interface mode. |
| P1(config-if)#ip address 11.1.1.31/24 | Assign the IPv4 address. |
| P1(config-if)#exit | Exit interface mode. |
| P1(config)#interface xe21 | Enter interface mode. |
| P1(config-if)#ip address 12.1.1.31/24 | Assign the IPv4 address. |
| P1(config-if)#exit | Exit interface mode. |
| P1(config)#router ospf 1 | Enter router OSPF mode. |
| P1(config-router)#ospf router-id 51.51.51.51 | Configure OSPF router id same as loopback ip address. |
| P1(config-router)#network 11.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| P1(config-router)#network 12.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| P1(config-router)#network 51.51.51.51/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| P1(config-router)#exit | Exit router OSPF mode. |
| P1(config-router)#router ldp | Enter router LDP mode. |
| P1(config-router)#router-id 51.51.51.51 | Set the router ID to IP address 51.51.51.51 |
| P1(config-router)#transport-address ipv4 51.51.51.51 | |
| P1(config-router)#exit | Exit LDP mode. |
| P1(config)#interface xe15 | Enter interface mode. |
| P1(config-if)#label-switching | Enable label switching in interface. |
| P1(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| P1(config-if)#exit | Exit interface mode. |
| P1(config)#interface xe21 | Enter interface mode. |
| P1(config-if)#label-switching | Enable label switching in interface. |
| P1(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| P1(config-if)#exit | Exit interface mode. |
| P1(config)#commit | Commit candidate configuration to be running configuration. |
| P1(config)#exit | Exit the config mode. |

PE2-ASBR1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| PE2-ASBR(config)#interface lo | Enter interface mode. |
| PE2-ASBR(config-if)#ip address 2.2.2.2/32 secondary | Assign the IPv4 address. |
| PE2-ASBR(config-if)#exit | Exit interface mode. |
| PE2-ASBR(config)#router ospf 1 | Enter router OSPF mode. |
| PE2-ASBR(config-router)#ospf router-id 2.2.2.2 | Configure OSPF router id same as loopback ip address. |
| PE2-ASBR(config-router)#network 2.2.2.2/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE2-ASBR(config-router)#network 12.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE2-ASBR(config-router)#exit | Exit router OSPF mode. |
| PE2-ASBR(config-router)#router ldp | Enter router LDP mode. |
| PE2-ASBR(config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| PE2-ASBR(config-router)#transport-address ipv4 2.2.2.2 | Configure LDP transport address same as loopback address. |
| PE2-ASBR(config-router)#exit | Exit LPD mode. |
| PE2-ASBR(config)#interface xe5 | Enter interface mode. |
| PE2-ASBR(config-if)#ip address 20.1.1.2/24 | Assign the IPv4 address. |
| PE2-ASBR(config-if)#label-switching | Enable label switching in interface. |
| PE2-ASBR(config-if)#exit | Exit interface mode. |
| PE2-ASBR(config)#interface xe20 | Enter interface mode. |
| PE2-ASBR(config-if)#ip address 12.1.1.2/24 | Assign the IPv4 address. |
| PE2-ASBR(config-if)#label-switching | Enable label switching in interface. |
| PE2-ASBR(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| PE2-ASBR(config-if)#exit | Exit interface mode |
| PE2-ASBR(config-if)#router bgp 65000 | Enter BGP router mode. |
| PE2-ASBR(config-router)# bgp router-id 2.2.2.2 | Configure BGP router-id. |
| PE2-ASBR(config-router)# no bgp default ipv4-unicast | Configure BGP to have no default ipv4 unicast capability exchange between neighbors. |
| PE2-ASBR(config-router)# neighbor 20.1.1.40 remote-as 65001 | Configure PE3-ASBR2 as an eBGP peer. |
| PE2-ASBR(config-router)# neighbor 41.41.41.41 remote-as 65000 | Configure PE1 as an iBGP peer. |
| PE2-ASBR(config-router)# neighbor 41.41.41.41 update-source lo | Update the source as loopback for iBGP peering with the remote PE1 router. |
| PE2-ASBR(config-router)# allocate-label all | Configure allocate-label. |
| PE2-ASBR(config-router)# address-family ipv4 | Enter address-family ipv4 labeled-unicast. |

| | |
|--|---|
| labeled-unicast | |
| PE2-ASBR(config-router-af)# neighbor 20.1.1.40 activate | Activate eBGP neighbor PE3-ASBR2. |
| PE2-ASBR(config-router-af)# neighbor 41.41.41.41 activate | Activate iBGP neighbor PE1. |
| PE2-ASBR(config-router-af)# neighbor 41.41.41.41 next-hop-self | Configure next-hop-self for iBGP neighbor PE1. |
| PE2-ASBR(config-router-af)# exit-address-family | Exit address-family ipv4 labeled-unicast. |
| PE2(config-router)#exit | Exit the router mode. |
| PE2(config)#commit | Commit candidate configuration to be running configuration. |
| PE2(config)#exit | Exit the config mode. |

PE3-ASBR

| | |
|---|--|
| PE3-ASBR#configure terminal | Enter configure mode. |
| PE3-ASBR(config)#interface lo | Enter interface mode. |
| PE3-ASBR(config-if)#ip address 40.40.40.40/32 secondary | Assign the IPv4 address. |
| PE3-ASBR(config-if)#exit | Exit interface mode. |
| PE3-ASBR(config)#interface xe5 | Enter interface mode. |
| PE3-ASBR(config-if)#ip address 20.1.1.40/24 | Assign the IPv4 address. |
| PE3-ASBR(config-if)#exit | Exit interface mode. |
| PE3-ASBR(config)#interface xe1 | Enter interface mode. |
| PE3-ASBR(config-if)#ip address 21.1.1.40/24 | Assign the IPv4 address. |
| PE3-ASBR(config-if)#exit | Exit interface mode. |
| PE3-ASBR(config)#router ospf 1 | Enter router OSPF mode. |
| PE3-ASBR(config-router)#ospf router-id 40.40.40.40 | Configure OSPF router id same as loopback ip address. |
| PE3-ASBR(config-router)#network 21.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE3-ASBR(config-router)#network 40.40.40.40/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE3-ASBR(config-router)#exit | Exit router OSPF mode. |
| PE3-ASBR(config-router)#router ldp | Enter router LDP mode. |
| PE3-ASBR(config-router)#router-id 40.40.40.40 | Set the router ID to IP address 40.40.40.40 |
| PE3-ASBR(config-router)#transport-address ipv4 40.40.40.40 | |
| PE3-ASBR(config)#interface xe1 | Enter interface mode. |
| PE3-ASBR(config-if)#label-switching | Enable label switching in interface. |
| PE3-ASBR(config-if)#enable-ldp ipv4 | Enable LDP on the interface. |

| | |
|--|---|
| PE3-ASBR(config)#interface xe5 | Enter interface mode. |
| PE3-ASBR(config-if)#label-switching | Enable label switching in interface. |
| PE3-ASBR(config-if)#router bgp 65001 | Enter BGP router mode. |
| PE3-ASBR(config-router)# bgp router-id 40.40.40.40 | Configure BGP router-id. |
| PE3-ASBR(config-router)# no bgp default ipv4-unicast | Configure BGP to have no default ipv4 unicast capability exchange with the neighbors. |
| PE3-ASBR(config-router)# neighbor 20.1.1.2 remote-as 65000 | Configure PE2-ASBR as an eBGP peer. |
| PE3-ASBR(config-router)# neighbor 28.28.28.28 remote-as 65001 | Configure PE4 as an iBGP peer. |
| PE3-ASBR(config-router)# neighbor 28.28.28.28 update-source lo | Update the source as loopback for iBGP peering with the remote PE4 router. |
| PE3-ASBR(config-router)# allocate-label all | Configure allocate-label. |
| PE3-ASBR(config-router)# address-family ipv4 labeled-unicast | Enter address-family ipv4 labeled-unicast. |
| PE3-ASBR(config-router-af)# neighbor 20.1.1.2 activate | Activate eBGP neighbor PE2-ASBR1. |
| PE3-ASBR(config-router-af)# neighbor 28.28.28.28 activate | Activate iBGP neighbor PE4. |
| PE3-ASBR(config-router-af)# neighbor 28.28.28.28 next-hop-self | Configure next-hop-self for iBGP neighbor PE4. |
| PE3-ASBR(config-router-af)# exit-address-family | Exit address-family ipv4 labeled-unicast. |
| PE3-ASBR(config-router)#exit | Exit the router mode. |
| PE3-ASBR(config)#commit | Commit candidate configuration to be running configuration. |
| PE3-ASBR(config)#exit | Exit the config mode. |

P2

| | |
|--|---|
| #configure terminal | Enter configure mode. |
| P2(config)#interface lo | Enter interface mode. |
| P2(config-if)# ip address 6.6.6.6/32 secondary | Assign the IPv4 address. |
| P2(config-if)#exit | Exit interface mode. |
| P2(config)#interface xe0 | Enter interface mode. |
| P2(config-if)#ip address 21.1.1.5/24 | Assign the IPv4 address. |
| P2(config-if)#exit | Exit interface mode. |
| P2(config)#interface xell | Enter interface mode. |
| P2(config-if)#ip address 22.1.1.5/24 | Assign the IPv4 address. |
| P2(config-if)#exit | Exit interface mode. |
| P2(config)#router ospf 1 | Enter router OSPF mode. |
| P2(config-router)#ospf router-id 6.6.6.6 | Configure OSPF router id same as loopback ip address. |

| | |
|--|--|
| P2(config-router)# network 6.6.6.6/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| P2(config-router)#network 21.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| P2(config-router)#network 22.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| P2(config-router)#exit | Exit router OSPF mode. |
| P2(config)#router ldp | Enter router LDP mode. |
| P2(config-router)#router-id 6.6.6.6 | Set the router ID to IP address 6.6.6.6 |
| P2(config-router)#transport-address ipv4 6.6.6.6 | Configure LDP transport address same as loopback address |
| P2(config-router)#exit | Exit LDP mode. |
| P2(config)#interface xe0 | Enter interface mode. |
| P2(config-if)#label-switching | Enable label switching in interface. |
| P2(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| P2(config-if)#exit | Exit interface mode. |
| P2(config)#interface xe11 | Enter interface mode. |
| P2(config-if)#label-switching | Enable label switching in interface. |
| P2(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| P2(config-if)#exit | Exit interface mode. |
| P2(config)#commit | Commit candidate configuration to be running configuration. |
| P2(config)#exit | Exit the config mode. |

PE4

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| PE4(config)#interface lo | Enter interface mode. |
| PE4(config-if)#ip address 28.28.28.28/32 secondary | Assign the IPv4 address. |
| PE4(config-if)#exit | Exit interface mode. |
| PE4(config)#ip vrf vrf1 | Create a new VRF named vrf1. |
| PE4(config-vrf)#rd 1:1 | Assign the route distinguisher (RD) value as 1:1. |
| PE4(config-vrf)#route-target both 1:1 | Configure import and export RTs for the VRF. |
| PE4(config-vrf)#exit | Exit VRF mode. |
| PE4(config)#interface xe23 | Enter interface mode. |
| PE4(config-if)#ip vrf forwarding vrf1 | Bind the interface connected to the CE router with VRF vrf1. |
| PE4(config-if)#ip address 30.1.1.1/24 | Assign the IPv4 address. |
| PE4(config-if)#exit | Exit interface mode. |

| | |
|---|--|
| PE4(config)#interface xe11 | Enter interface mode. |
| PE4(config-if)#ip address 22.1.1.28/24 | Assign the IPv4 address. |
| PE4(config-if)#exit | Exit interface mode. |
| PE4(config)#router ospf 1 | Enter router OSPF mode. |
| PE4(config-router)#ospf router-id 28.28.28.28 | Configure OSPF router id same as loopback ip address. |
| PE4(config-router)#network 22.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE4(config-router)#network 28.28.28.28/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| PE4(config-router)#exit | Exit OSPF router mode. |
| PE4(config)# router ldp | Enter router LDP mode. |
| PE4(config-router)#router-id 28.28.28.28 | Set the router ID to IP address 28.28.28.28 |
| PE4(config-router)#transport-address ipv4 28.28.28.28 | Configure LDP transport address same as loopback address |
| PE4(config-router)#exit | Exit LDP mode. |
| PE4(config)#interface xe11 | Enter interface mode. |
| PE4(config-if)#label-switching | Enable label switching in interface. |
| PE4(config-if)# enable-ldp ipv4 | Enable LDP on the interface. |
| PE4(config-if)#exit | Exit interface mode. |
| PE4(config-if)#router bgp 65001 | Enter BGP router mode. |
| PE4(config-router)# bgp router-id 28.28.28.28 | Configure BGP router-id. |
| PE4(config-router)# no bgp default ipv4-unicast | Configure BGP speaker to have no default ipv4 unicast capability exchange between neighbors. |
| PE4(config-router)#address-family ipv4 unicast | Enter the IPv4 unicast address family. |
| PE4(config-router-af)#network 28.28.28.28/32 | Advertise loopback address. |
| PE4(config-router-af)#exit-address-family | Exit address family mode |
| PE4(config-router)# neighbor 40.40.40.40 remote-as 65001 | Configure PE3-ASBR2 as an iBGP peer. |
| PE4(config-router)# neighbor 40.40.40.40 update-source lo | Update the source as loopback for iBGP peering with the remote PE3-ASBR router. |
| PE4(config-router)# neighbor 41.41.41.41 remote-as 65000 | Configure PE1 as an eBGP peer. |
| PE4(config-router)# neighbor 41.41.41.41 ebgp-multihop | Configure eBGP multichip for eBGP peer PE1. |
| PE4(config-router)# neighbor 41.41.41.41 update-source lo | Update the source as loopback for eBGP peering with the remote PE1 router. |
| PE4(config-router)# allocate-label all | Configure allocate-label. |
| PE4(config-router)# address-family ipv4 la-beled-unicast | Enter address-family ipv4 labeled-unicast. |

| | |
|--|---|
| PE4(config-router-af)# neighbor 40.40.40.40 activate | Activate iBGP neighbor. |
| PE4(config-router-af)# exit-address-family | Exit address-family ipv4 labeled-unicast. |
| PE4(config-router)# address-family vpnv4 unicast | Enter address-family vpnv4. |
| PE4(config-router-af)# neighbor 41.41.41.41 allow-ebgp-vpn | Activate eBGP neighbor PE1. |
| PE4(config-router-af)# neighbor 41.41.41.41 activate | Configure allow-ebgp-vpn for ebgp neighbor PE1. |
| PE4(config-router-af)# exit-address-family | Exit address-family vpnv4. |
| PE4(config-router)# address-family ipv4 vrf vrf1 | Enter the IPv4 address family for VRF vrf1. |
| PE4(config-router-af)# redistribute con-nected | Redistribute connected route. |
| PE4(config-router-af)# exit-address-family | Exit IPv4 VRF Address Family mode. |
| PE4(config-router)#exit | Exit the router mode. |
| PE4(config)#commit | Commit candidate configuration to be running configuration. |
| PE4(config)#exit | Exit the config mode. |

Validation

PE1

```

PE1#show ip bgp labeled-unicast summary
BGP router identifier 41.41.41.41, local AS number 65000
BGP table version is 10
2 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd
2.2.2.2            4 65000   374      368      10      0      0   02:37:43      3

Total number of neighbors 1

Total number of Established sessions 1

PE1#show ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale
   Network          Next Hop        In Label    Out Label
*>i 5.5.5.5/32       2.2.2.2          24960       24965
*>i 28.28.28.28/32   2.2.2.2          24962       24962
*>i 31.31.31.31/32   2.2.2.2          24963       24961
*>  41.41.41.41/32   0.0.0.0          24961       -

PE1#show ip route bgp
IP Route Table for VRF "default"
B          28.28.28.28/32 [200/0] via 2.2.2.2 (recursive via 11.1.1.31), 02:41:00

Gateway of last resort is not set

PE1#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code    FEC                FTN-ID   Tunnel-id  Pri   LSP-Type        Out-Label  Out-
Intf    ELC                Nexthop

```



```

L> 2.2.2.2/32 4 0 Yes LSP_
DEFAULT 24321 xe15 No 11.1.1.31
L> 5.5.5.5/32 6 0 Yes LSP_
DEFAULT 24324 xe15 No 11.1.1.31
B 5.5.5.5/32 7 0 Yes LSP_DEFAULT 24965 -
No 2.2.2.2
L> 12.1.1.0/24 2 0 Yes LSP_
DEFAULT 3 xe15 No 11.1.1.31
L> 13.1.1.0/24 3 0 Yes LSP_
DEFAULT 3 xe15 No 11.1.1.31
B> 28.28.28.28/32 8 0 Yes LSP_DEFAULT 24962 -
No 2.2.2.2
L> 31.31.31.31/32 5 0 Yes LSP_
DEFAULT 24322 xe15 No 11.1.1.31
B 31.31.31.31/32 9 0 Yes LSP_DEFAULT 24961 -
No 2.2.2.2
L> 51.51.51.51/32 1 0 Yes LSP_
DEFAULT 3 xe15 No 11.1.1.31

PE1#show ip bgp vpnv4 all summary
BGP router identifier 41.41.41.41, local AS number 65000
BGP table version is 4
2 BGP AS-PATH entries
0 BGP community entries

Neighbor V AS MsgRcv MsgSen TblVer InQ OutQ Up/Down State/PfxRcd
28.28.28.28 4 65001 340 338 4 0 0 02:22:11 1

Total number of neighbors 1

Total number of Established sessions 1

PE1#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path
Route Distinguisher: 1:1 (Default for VRF vrf1)
*> 1 10.1.1.0/24 0.0.0.0 0 100 32768 ?
*> 30.1.1.0/24 28.28.28.28 0 100 0 65001 ?
Announced routes count = 1
Accepted routes count = 1
Route Distinguisher: 1:1
*> 30.1.1.0/24 28.28.28.28 0 100 0 65001 ?
Announced routes count = 0
Accepted routes count = 1

PE1#show ip route vrf vrf1 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN,
v - vrf leaked
> - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1"
C *> 10.1.1.0/24 is directly connected, xe23, 1d19h57m
B *> 30.1.1.0/24 [20/0] via 28.28.28.28, 02:22:28

Gateway of last resort is not set

```

P1

```

P1#show mpls ilm-table

```

```

Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, S - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|----------------|--------|----------|-----------|---------|---------------|
| Intf/VRF | Nexthop | | LSP-Type | | | |
| L> | 2.2.2.2/32 | 2 | 24321 | 3 | N/A | xe21 12.1.1.2 |
| | LSP_DEFAULT | | | | | |
| L> | 41.41.41.41/32 | 1 | 24320 | 3 | N/A | xe15 11.1.1.2 |
| | LSP_DEFAULT | | | | | |
| L> | 31.31.31.31/32 | 4 | 24323 | Nolabel | N/A | N/A 127.0.0.1 |
| | LSP_DEFAULT | | | | | |
| L> | 31.31.31.31/32 | 3 | 24322 | Nolabel | N/A | N/A 127.0.0.1 |
| | LSP_DEFAULT | | | | | |
| L> | 5.5.5.5/32 | 5 | 24324 | Nolabel | N/A | N/A 127.0.0.1 |
| | LSP_DEFAULT | | | | | |

PE2-ASBR

```

PE2-ASBR#show ip bgp labeled-unicast summary
BGP router identifier 2.2.2.2, local AS number 65000
BGP table version is 10
2 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|-------------|---|-------|--------|--------|--------|-----|------|----------|--------------|
| 20.1.1.40 | 4 | 65001 | 6427 | 6521 | 10 | 0 | 0 | 1d20h28m | 2 |
| 41.41.41.41 | 4 | 65000 | 379 | 385 | 10 | 0 | 0 | 02:42:38 | 1 |

Total number of neighbors 2

Total number of Established sessions 2

```
PE2-ASBR#show mpls ilm-table
```

```

Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, S - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|----------------|--------|----------|-----------|---------|----------------|
| Intf/VRF | Nexthop | | LSP-Type | | | |
| B> | 31.31.31.31/32 | 2 | 24961 | Nolabel | N/A | N/A 127.0.0.1 |
| | LSP_DEFAULT | | | | | |
| B> | 41.41.41.41/32 | 1 | 24960 | 24961 | N/A | N/A 41.41.41.4 |
| | LSP_DEFAULT | | | | | |
| B> | 5.5.5.5/32 | 6 | 24965 | 24965 | N/A | N/A 20.1.1.40 |
| | LSP_DEFAULT | | | | | |
| B> | 28.28.28.28/32 | 3 | 24962 | 24961 | N/A | N/A 20.1.1.40 |
| | LSP_DEFAULT | | | | | |

```
PE2-ASBR#show mpls forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

```

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|-------------|---------|-----------|-----------|----------|-----------|------|
| Intf | ELC | Nexthop | | | | | |
| B> | 5.5.5.5/32 | 8 | 0 | Yes | LSP_ | | |
| DEFAULT | 24965 | xe5 | No | 20.1.1.40 | | | |
| L> | 11.1.1.0/24 | 1 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe20 | No | 12.1.1.31 | | | |
| L> | 13.1.1.0/24 | 2 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe20 | No | 12.1.1.31 | | | |

```

B> 28.28.28.28/32 7 0 Yes LSP_
DEFAULT 24961 xe5 No 20.1.1.40
L> 31.31.31.31/32 5 0 Yes LSP_
DEFAULT 24323 xe20 No 12.1.1.31
L> 41.41.41.41/32 3 0 Yes LSP_
DEFAULT 24320 xe20 No 12.1.1.31
B 41.41.41.41/32 6 0 Yes LSP_DEFAULT 24961 -
No 41.41.41.41
L> 51.51.51.51/32 4 0 Yes LSP_
DEFAULT 3 xe20 No 12.1.1.31

PE2-ASBR#show ip route bgp
IP Route Table for VRF "default"
B 28.28.28.28/32 [20/0] via 20.1.1.40, xe5, 02:41:38

Gateway of last resort is not set

```

PE3-ASBR

```

PE3-ASBR#show ip bgp labeled-unicast summary
BGP router identifier 40.40.40.40, local AS number 65001
BGP table version is 11
2 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd
20.1.1.2           4 65000  6263     6267    11       0      0   1d20h31m         2
28.28.28.28        4 65001   383      389    11       0      0   02:42:56         1

Total number of neighbors 2

Total number of Established sessions 2

PE3-ASBR#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, S - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

Code  FEC/VRF/L2CKT  ILM-ID  In-Label  Out-Label  In-Intf  Out-
Intf/VRF  Nexthop      LSP-Type
B> 31.31.31.31/32  6        24964     24961      N/A       N/A       20.1.1.2
      LSP_DEFAULT
B> 41.41.41.41/32  2        24960     24960      N/A       N/A       20.1.1.2
      LSP_DEFAULT
B> 28.28.28.28/32  3        24961     24321      N/A       N/A       28.28.28.2
8      LSP_DEFAULT
B> 5.5.5.5/32      7        24965     Nolabel    N/A       N/A       127.0.0.1
      LSP_DEFAULT

PE3-ASBR#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
       B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

Code  FEC          FTN-ID  Tunnel-id  Pri  LSP-Type  Out-Label  Out-
Intf  ELC      Nexthop
L> 5.5.5.5/32  5        0        Yes  LSP_
DEFAULT 24320  xe1      No    21.1.1.5
L> 6.6.6.6/32  6        0        Yes  LSP_
DEFAULT 3      xe1      No    21.1.1.5
L> 22.1.1.0/24 7        0        Yes  LSP_
DEFAULT 3      xe1      No    21.1.1.5
L> 23.1.1.0/24 8        0        Yes  LSP_
DEFAULT 3      xe1      No    21.1.1.5
L> 28.28.28.28/32 9        0        Yes  LSP_

```

```

DEFAULT      24321      xe1      No      21.1.1.5
  B      28.28.28.28/32      2      0      Yes      LSP_DEFAULT      24321      -
           No      28.28.28.28
  B>      31.31.31.31/32      4      0      Yes      LSP_
DEFAULT      24961      xe5      No      20.1.1.2
  B>      41.41.41.41/32      1      0      Yes      LSP_
DEFAULT      24960      xe5      No      20.1.1.2

```

```

PE3-ASBR#show ip route bgp
IP Route Table for VRF "default"
B      41.41.41.41/32 [20/0] via 20.1.1.2, xe5, 02:45:37

```

Gateway of last resort is not set

P2

```

P2#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, S - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- | |
|----------|----------------|--------|----------|-----------|---------|------|-----------|
| Intf/VRF | NextHop | | LSP-Type | | | | |
| L> | 5.5.5.5/32 | 3 | 24322 | Nolabel | N/A | N/A | 127.0.0.1 |
| | LSP_DEFAULT | | | | | | |
| L> | 5.5.5.5/32 | 1 | 24320 | Nolabel | N/A | N/A | 127.0.0.1 |
| | LSP_DEFAULT | | | | | | |
| L> | 28.28.28.28/32 | 2 | 24321 | 3 | N/A | xe11 | 22.1.1.28 |
| | LSP_DEFAULT | | | | | | |
| L> | 31.31.31.31/32 | 4 | 24323 | Nolabel | N/A | N/A | 127.0.0.1 |
| | LSP_DEFAULT | | | | | | |
| L> | 40.40.40.40/32 | 5 | 24324 | 3 | N/A | xe0 | 21.1.1.40 |
| | LSP_DEFAULT | | | | | | |

PE4

```

PE4#show ip bgp labeled-unicast summary
BGP router identifier 28.28.28.28, local AS number 65001
BGP table version is 12
2 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|-------------|---|-------|--------|--------|--------|-----|------|----------|--------------|
| 40.40.40.40 | 4 | 65001 | 404 | 399 | 12 | 0 | 0 | 02:49:41 | 3 |

Total number of neighbors 1

Total number of Established sessions 1

```
PE4#show ip bgp labeled-unicast
```

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, S - stale

| Network | Next Hop | In Label | Out Label |
|--------------------|-------------|----------|-----------|
| *>i 5.5.5.5/32 | 40.40.40.40 | 24324 | 24965 |
| *> 28.28.28.28/32 | 0.0.0.0 | 24321 | - |
| *>i 31.31.31.31/32 | 40.40.40.40 | 24322 | 24964 |
| *>i 41.41.41.41/32 | 40.40.40.40 | 24323 | 24960 |

PE4#

```
PE4#show mpls forwarding-table
```

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,

L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|----------------|-------------|-----------|----------|-------------|-----------|------|
| Intf | ELC | Nextthop | | | | | |
| L> | 5.5.5.5/32 | 4 | 0 | Yes | LSP_ | | |
| DEFAULT | 24322 | xe11 | No | 22.1.1.5 | | | |
| B | 5.5.5.5/32 | 3 | 0 | Yes | LSP_DEFAULT | 24965 | - |
| | No | 40.40.40.40 | | | | | |
| L> | 6.6.6.6/32 | 5 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe11 | No | 22.1.1.5 | | | |
| L> | 21.1.1.0/24 | 6 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe11 | No | 22.1.1.5 | | | |
| L> | 23.1.1.0/24 | 7 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe11 | No | 22.1.1.5 | | | |
| L> | 31.31.31.31/32 | 8 | 0 | Yes | LSP_ | | |
| DEFAULT | 24323 | xe11 | No | 22.1.1.5 | | | |
| B | 31.31.31.31/32 | 1 | 0 | Yes | LSP_DEFAULT | 24964 | - |
| | No | 40.40.40.40 | | | | | |
| L> | 40.40.40.40/32 | 9 | 0 | Yes | LSP_ | | |
| DEFAULT | 24324 | xe11 | No | 22.1.1.5 | | | |
| B> | 41.41.41.41/32 | 2 | 0 | Yes | LSP_DEFAULT | 24960 | - |
| | No | 40.40.40.40 | | | | | |

```
PE4#show ip bgp vpnv4 all summary
BGP router identifier 28.28.28.28, local AS number 65001
BGP table version is 4
2 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|-------------|---|-------|--------|--------|--------|-----|------|----------|--------------|
| 41.41.41.41 | 4 | 65000 | 373 | 376 | 4 | 0 | 0 | 02:37:04 | 1 |

Total number of neighbors 1

Total number of Established sessions 1

```
PE4#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|---|-------------|--------|--------|--------|---------|
| Route Distinguisher: 1:1 (Default for VRF vrf1) | | | | | |
| *> 10.1.1.0/24 | 41.41.41.41 | 0 | 100 | 0 | 65000 ? |
| *> 1 30.1.1.0/24 | 0.0.0.0 | 0 | 100 | 32768 | ? |
| Announced routes count = 1 | | | | | |
| Accepted routes count = 1 | | | | | |
| Route Distinguisher: 1:1 | | | | | |
| *> 10.1.1.0/24 | 41.41.41.41 | 0 | 100 | 0 | 65000 ? |
| Announced routes count = 0 | | | | | |
| Accepted routes count = 1 | | | | | |

```
PE4#show ip route vrf vrf1 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN,
v - vrf leaked
> - selected route, * - FIB route, p - stale info
```

```
IP Route Table for VRF "vrf1"
B *> 10.1.1.0/24 [20/0] via 41.41.41.41, 02:37:23
C *> 30.1.1.0/24 is directly connected, xe21, 1d20h11m
```

Gateway of last resort is not set

```

PE4#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, S - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|----------------|--------|----------|-----------|---------|------------|
| Intf/VRF | NextHop | | LSP-Type | | | |
| B> | 41.41.41.41/32 | 4 | 24323 | 24960 | N/A | N/A |
| 0 | LSP_DEFAULT | | | | | 40.40.40.4 |
| B> | 28.28.28.28/32 | 2 | 24321 | Nolabel | N/A | N/A |
| | LSP_DEFAULT | | | | | 127.0.0.1 |
| B> | vrf1 | 1 | 24320 | Nolabel | N/A | vrf1 |
| | LSP_DEFAULT | | | | | N/A |
| B> | 31.31.31.31/32 | 3 | 24322 | 24964 | N/A | N/A |
| 0 | LSP_DEFAULT | | | | | 40.40.40.4 |
| B> | 5.5.5.5/32 | 5 | 24324 | 24965 | N/A | N/A |
| 0 | LSP_DEFAULT | | | | | 40.40.40.4 |

ECMP Support for L3VPN

Overview


Equal Cost Multipath (ECMP) support for L3VPN helps in load-balancing the traffic for L3VPN routes while managing traffic distribution across multiple paths of equal cost. It enables improved load balancing and resiliency across VPN traffic paths.

Currently, in Inter-AS VPN configuration options A, B, and C, the best path is programmed in fast path for L3VPN routes. Using ECMP, multiple paths are configured in fast path while also balancing the traffic load for a given L3VPN route.

Feature Characteristics

The PE routers PE1, PE2, ... P8 learn IPv4 and IPv6 routes connected to VRF on the User Network Interface. The protocol BGP running in the devices publish these routes to UPE as VPNv4/VPNv6 routes using Multi Protocol BGP with Route Distinguisher (RD) and Route Target (RT). The user facing PE installs the next-hops to its Routing Information Base (RIB) and Forwarding Information Base (FIB) and prepares the hashing keys based on the L3 fields of the incoming traffic, and also load balances the traffic at service level and transport (LDP/RSVP/SR) level.

Figure 93. L3VPN ECMP Characteristics



The Hardware Service Layer (HSL) programs the MPLS forwarding entities in hardware, maintains the software linking of the programmed entities and is responsible for load balancing of L3VPN traffic. It resolves Forward Error Correction (FEC) and uses the hierarchical FEC resolution along with the ECMP group table. The existing L3 keys provided by OcNOS are used to load balance the traffic, and MPLS Label stack is used to load balance the traffic at transit nodes.

Load balancing is done at service and transport level and multiple paths are maintained at L3VPN for the given nexthops. The first load balance is done at service level hierarchy and the second load balance is done at transport level hierarchy, such as LDP, RSVP, and SR.

The existing load balance keys are used at the PE ingress nodes and the incoming label stack is used at the P transit nodes (ASBR and RR) to load balance the traffic.

Supported Hardware

Qumran1, Qumran2, and J2C+ devices.

Limitations

- Load balancing of transit routers (RR and ASBRs) in case of inter-AS is not supported.
- Multipath EVPN Integrated Routing and Bridging (IRB) is not supported.

Prerequisites

1. Define Interfaces and Loopback Addresses

Configure Layer 2 interfaces such as port channel interfaces, for example po1, and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity.

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 8.8.8.8/32 secondary  
  ipv6 address ::1/128  
  
interface po7  
  ip address 31.1.1.8/24
```

2. Configure IGP for Dynamic Routing

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

ISIS Configuration

```
router isis 1  
  is-type level-2-only  
  metric-style wide  
  microloop-avoidance level-2  
  mpls traffic-eng router-id 8.8.8.8  
  mpls traffic-eng level-2  
  capability cspf  
  dynamic-hostname  
  bfd all-interfaces  
  net 49.0000.0000.0008.00  
  passive-interface lo  
!  
interface po7  
  mpls ldp-igp sync isis level-2
```

```
isis network point-to-point
ip router isis 1
```

OSPF Configuration

```
router ospf 1
ospf router-id 8.8.8.8
network 8.8.8.8/32 area 0.0.0.0
network 31.1.1.0/24 area 0.0.0.0!
!
interface po7
ip ospf network point-to-point
```

Configuration


To set up ECMP for L3VPN, follow the configuration steps mentioned below.

Topology

In this topology, the PE1 router installs the next-hops to its routing and forwarding information base (RIB and FIB). It prepares the hashing keys based on the L3 fields of the incoming traffic and load balances the traffic. The same IPv4 (200.1.1.0/24)/ IPv6(2000::/64) prefixes are advertised via BGP using L3VPN from PE2 and PE3 routers.

PE1 installs multiple next hops to these prefixes in the VPN FIB.

Figure 94. L3VPN ECMP Configuration



PE1

1. Configure Router LDP

```
PE1(config)#router ldp
PE1(config-router)# router-id 8.8.8.8
PE1(config-router)# transport-address ipv4 8.8.8.8
PE1(config-router)#
```

2. Enable LDP and label-switching for core interface

```
PE1(config)#interface po7
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching

PE1(config)#interface po10
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching
```

3. Configure VRF

```
PE1(config)#ip vrf vrf1000
```



```
PE1(config-vrf)# rd 90000:8
PE1(config-vrf)# route-target both 90000:1
PE1(config-vrf)#
```

4. Configure subinterface and attach VRF to subinterface

```
PE1(config)#interface xe26.100
PE1(config-if)# encapsulation dot1q 100
PE1(config-if)# ip vrf forwarding vrf1000
PE1(config-if)# ip address 100.1.1.1/24
PE1(config-if)# ipv6 address 1000::1/24
```

5. Configure BGP

```
PE1(config)#router bgp 90000
PE1(config-router)# bgp router-id 8.8.8.8
PE1(config-router)# neighbor 4.4.4.4 remote-as 90000
PE1(config-router)# neighbor 4.4.4.4 update-source lo
PE1(config-router)# neighbor 7.7.7.7 remote-as 90000
PE1(config-router)# neighbor 7.7.7.7 update-source lo
PE1(config-router)# !
PE1(config-router)# address-family vpnv4 unicast
PE1(config-router-af)# neighbor 4.4.4.4 activate
PE1(config-router-af)# neighbor 7.7.7.7 activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# address-family vpnv6 unicast
PE1(config-router-af)# neighbor 4.4.4.4 activate
PE1(config-router-af)# neighbor 7.7.7.7 activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# address-family ipv4 vrf vrf1000
PE1(config-router-af)# redistribute connected
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# address-family ipv6 vrf vrf1000
PE1(config-router-af)# redistribute connected
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# exit
PE1(config)#
```

6. Configure Multipath and max-paths iBGP

```
PE1(config)#router bgp 90000
PE1(config-router)#address-family vpnv4
PE1(config-router-af)#multi-path
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family vpnv6
PE1(config-router-af)#multi-path
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family ipv4 vrf vrf1000
PE1(config-router-af)#ib
PE1(config-router-af)#ib
PE1(config-router-af)#max-paths ibgp 10
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family ipv6 vrf vrf1000
PE1(config-router-af)#max-paths ibgp 10
PE1(config-router-af)#
PE1(config-router-af)#exit
PE1(config-router)#
```

RR1

1. Configure Router LDP

```
RR1(config)#router ldp
```

```
RR1(config-router)# router-id 4.4.4.4
RR1(config-router)# transport-address ipv4 4.4.4.4
RR1(config-router)#
```

2. Enable LDP and label-switching for core interface

```
RR1(config)#interface po7
RR1(config-if)# enable-ldp ipv4
RR1(config-if)#label-switching

RR1(config)#interface sa4
RR1(config-if)# enable-ldp ipv4
RR1(config-if)#label-switching

RR1(config)#interface xe26
RR1(config-if)# enable-ldp ipv4
RR1(config-if)#label-switching
```

3. Configure BGP

```
RR1(config)#router bgp 90000
RR1(config-router)# bgp router-id 4.4.4.4
RR1(config-router)# no bgp inbound-route-filter
RR1(config-router)# neighbor 3.3.3.3 remote-as 90000
RR1(config-router)# neighbor 3.3.3.3 update-source lo
RR1(config-router)# neighbor 5.5.5.5 remote-as 90000
RR1(config-router)# neighbor 5.5.5.5 update-source lo
RR1(config-router)# neighbor 8.8.8.8 remote-as 90000
RR1(config-router)# neighbor 8.8.8.8 update-source lo
RR1(config-router)# !
RR1(config-router)# address-family vpnv4 unicast
RR1(config-router-af)# neighbor 3.3.3.3 activate
RR1(config-router-af)# neighbor 3.3.3.3 route-reflector-client
RR1(config-router-af)# neighbor 3.3.3.3 next-hop-self
RR1(config-router-af)# neighbor 5.5.5.5 activate
RR1(config-router-af)# neighbor 5.5.5.5 route-reflector-client
RR1(config-router-af)# neighbor 5.5.5.5 next-hop-self
RR1(config-router-af)# neighbor 8.8.8.8 activate
RR1(config-router-af)# neighbor 8.8.8.8 route-reflector-client
RR1(config-router-af)# neighbor 8.8.8.8 next-hop-self
RR1(config-router-af)# exit-address-family
RR1(config-router)# !
RR1(config-router)# address-family vpnv6 unicast
RR1(config-router-af)# multi-path
RR1(config-router-af)# neighbor 3.3.3.3 activate
RR1(config-router-af)# neighbor 3.3.3.3 route-reflector-client
RR1(config-router-af)# neighbor 3.3.3.3 next-hop-self
RR1(config-router-af)# neighbor 5.5.5.5 activate
RR1(config-router-af)# neighbor 5.5.5.5 route-reflector-client
RR1(config-router-af)# neighbor 5.5.5.5 next-hop-self
RR1(config-router-af)# neighbor 8.8.8.8 activate
RR1(config-router-af)# neighbor 8.8.8.8 route-reflector-client
RR1(config-router-af)# neighbor 8.8.8.8 next-hop-self
RR1(config-router-af)# exit-address-family
RR1(config-router)#
```

RR2

1. Configure Router LDP

```
RR2(config)#router ldp
RR2(config-router)# router-id 7.7.7.7
RR2(config-router)# transport-address ipv4 7.7.7.7
RR2(config-router)#
```

2. Enable LDP and label-switching for core interface

```
RR2(config)#interface po75
RR2(config-if)# enable-ldp ipv4
RR2(config-if)#label-switching

RR2(config)#interface po10
RR2(config-if)# enable-ldp ipv4
RR2(config-if)#label-switching

RR2(config)#interface xe6
RR2(config-if)# enable-ldp ipv4
RR2(config-if)#label-switching
```

3. Configure BGP

```
RR2(config)#router bgp 90000
RR2(config-router)# bgp router-id 7.7.7.7
RR2(config-router)# no bgp inbound-route-filter
RR2(config-router)# neighbor 3.3.3.3 remote-as 90000
RR2(config-router)# neighbor 3.3.3.3 update-source lo
RR2(config-router)# neighbor 5.5.5.5 remote-as 90000
RR2(config-router)# neighbor 5.5.5.5 update-source lo
RR2(config-router)# neighbor 8.8.8.8 remote-as 90000
RR2(config-router)# neighbor 8.8.8.8 update-source lo
RR2(config-router)# !
RR2(config-router)# address-family vpnv4 unicast
RR2(config-router-af)# neighbor 3.3.3.3 activate
RR2(config-router-af)# neighbor 3.3.3.3 route-reflector-client
RR2(config-router-af)# neighbor 3.3.3.3 next-hop-self
RR2(config-router-af)# neighbor 5.5.5.5 activate
RR2(config-router-af)# neighbor 5.5.5.5 route-reflector-client
RR2(config-router-af)# neighbor 5.5.5.5 next-hop-self
RR2(config-router-af)# neighbor 8.8.8.8 activate
RR2(config-router-af)# neighbor 8.8.8.8 route-reflector-client
RR2(config-router-af)# neighbor 8.8.8.8 next-hop-self
RR2(config-router-af)# exit-address-family
RR2(config-router)# !
RR2(config-router)# address-family vpnv6 unicast
RR2(config-router-af)# neighbor 3.3.3.3 activate
RR2(config-router-af)# neighbor 3.3.3.3 route-reflector-client
RR2(config-router-af)# neighbor 3.3.3.3 next-hop-self
RR2(config-router-af)# neighbor 5.5.5.5 activate
RR2(config-router-af)# neighbor 5.5.5.5 route-reflector-client
RR2(config-router-af)# neighbor 5.5.5.5 next-hop-self
RR2(config-router-af)# neighbor 8.8.8.8 activate
RR2(config-router-af)# neighbor 8.8.8.8 route-reflector-client
RR2(config-router-af)# neighbor 8.8.8.8 next-hop-self
RR2(config-router-af)# exit-address-family
RR2(config-router)#
```

PE2

1. Configure Router LDP

```
PE2(config)#router ldp
PE2(config-router)# router-id 5.5.5.5
PE2(config-router)# transport-address ipv4 5.5.5.5
PE2(config-router)#
```

2. Enable LDP and label-switching for core interface

```
PE2(config)#interface xe0
PE2(config-if)# enable-ldp ipv4
PE2(config-if)#label-switching

PE2(config)#interface po75
PE2(config-if)# enable-ldp ipv4
PE2(config-if)#label-switching
```

3. Configure VRF

```
PE2(config)#ip vrf vrf1000
PE2(config-vrf)# rd 90000:5
PE2(config-vrf)# route-target both 90000:1
PE2(config-vrf)#
```

4. Configure sub-interface and attach the VRF to sub-interface

```
PE2(config)#interface xe2.100
PE2(config-if)# encapsulation dot1q 100
PE2(config-if)# ip vrf forwarding vrf1000
PE2(config-if)# ip address 200.1.1.5/24
PE2(config-if)# ipv6 address 2000::5/64
PE2(config-if)#
```

5. Configure BGP

```
PE2(config)#router bgp 90000
PE2(config-router)# bgp router-id 5.5.5.5
PE2(config-router)# neighbor 4.4.4.4 remote-as 90000
PE2(config-router)# neighbor 4.4.4.4 update-source lo
PE2(config-router)# neighbor 7.7.7.7 remote-as 90000
PE2(config-router)# neighbor 7.7.7.7 update-source lo
PE2(config-router)# !
PE2(config-router)# address-family vpnv4 unicast
PE2(config-router-af)# neighbor 4.4.4.4 activate
PE2(config-router-af)# neighbor 7.7.7.7 activate
PE2(config-router-af)# exit-address-family
PE2(config-router)# !
PE2(config-router)# address-family vpnv6 unicast
PE2(config-router-af)# neighbor 4.4.4.4 activate
PE2(config-router-af)# neighbor 7.7.7.7 activate
PE2(config-router-af)# exit-address-family
PE2(config-router)# !
PE2(config-router)# address-family ipv4 vrf vrf1000
PE2(config-router-af)# redistribute connected
PE2(config-router-af)# exit-address-family
PE2(config-router)# !
PE2(config-router)# address-family ipv6 vrf vrf1000
PE2(config-router-af)# redistribute connected
PE2(config-router-af)# exit-address-family
PE2(config-router)#
```

PE3

1. Configure Router LDP

```
PE3(config)#router ldp
PE3(config-router)# router-id 3.3.3.3
PE3(config-router)# transport-address ipv4 3.3.3.3
PE3(config-router)#
```

2. Enable LDP and label-switching for core interface

```
PE3(config)#interface xe6
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching

PE3(config)#interface sa4
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching
```

3. Configure VRF

```
PE3(config)#ip vrf vrf1000
PE3(config-vrf)# rd 90000:3
PE3(config-vrf)# route-target both 90000:1
PE3(config-vrf)#
```

4. Configure sub-interface and attach the VRF to sub-interface

```
PE3(config)#interface xe6.100
PE3(config-if)# encapsulation dot1q 100
PE3(config-if)# ip vrf forwarding vrf1000
PE3(config-if)# ip address 200.1.1.3/24
PE3(config-if)# ipv6 address 2000::3/64
PE3(config-if)#
```

5. Configure BGP

```
PE3(config)#router bgp 90000
PE3(config-router)# bgp router-id 3.3.3.3
PE3(config-router)# neighbor 4.4.4.4 remote-as 90000
PE3(config-router)# neighbor 4.4.4.4 update-source lo
PE3(config-router)# neighbor 7.7.7.7 remote-as 90000
PE3(config-router)# neighbor 7.7.7.7 update-source lo
PE3(config-router)# !
PE3(config-router)# address-family vpnv4 unicast
PE3(config-router-af)# neighbor 4.4.4.4 activate
PE3(config-router-af)# neighbor 7.7.7.7 activate
PE3(config-router-af)# exit-address-family
PE3(config-router)# !
PE3(config-router)# address-family vpnv6 unicast
PE3(config-router-af)# neighbor 4.4.4.4 activate
PE3(config-router-af)# neighbor 7.7.7.7 activate
PE3(config-router-af)# exit-address-family
PE3(config-router)# !
PE3(config-router)# address-family ipv4 vrf vrf1000
PE3(config-router-af)# redistribute connected
PE3(config-router-af)# exit-address-family
PE3(config-router)# !
PE3(config-router)# address-family ipv6 vrf vrf1000
PE3(config-router-af)# redistribute connected
PE3(config-router-af)# exit-address-family
PE3(config-router)#
```

Configuration Snapshot

PE1

VRF

```
ip vrf vrf1000
rd 90000:8
route-target both 90000:1
```

BGP

```
bgp extended-asn-cap
!
router bgp 90000
  bgp router-id 8.8.8.8
  neighbor 4.4.4.4 remote-as 90000
  neighbor 4.4.4.4 update-source lo
  neighbor 7.7.7.7 remote-as 90000
  neighbor 7.7.7.7 update-source lo
  !
  address-family vpnv4 unicast
  multi-path
  neighbor 4.4.4.4 activate
  neighbor 7.7.7.7 activate
  exit-address-family
  !
  address-family vpnv6 unicast
  multi-path
  neighbor 4.4.4.4 activate
```

```
neighbor 7.7.7.7 activate
exit-address-family
!
address-family ipv4 vrf vrf1000
max-paths ibgp 10
redistribute connected
exit-address-family
!
address-family ipv6 vrf vrf1000
max-paths ibgp 10
redistribute connected
exit-address-family
!
Exit
```

RR1

BGP

```
bgp extended-asn-cap
!
router bgp 90000
  bgp router-id 4.4.4.4
  no bgp inbound-route-filter
  neighbor 3.3.3.3 remote-as 90000
  neighbor 3.3.3.3 update-source lo
  neighbor 5.5.5.5 remote-as 90000
  neighbor 5.5.5.5 update-source lo
  neighbor 8.8.8.8 remote-as 90000
  neighbor 8.8.8.8 update-source lo
  !
  address-family vpnv4 unicast
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 route-reflector-client
  neighbor 3.3.3.3 next-hop-self
  neighbor 5.5.5.5 activate
  neighbor 5.5.5.5 route-reflector-client
  neighbor 5.5.5.5 next-hop-self
  neighbor 8.8.8.8 activate
  neighbor 8.8.8.8 route-reflector-client
  neighbor 8.8.8.8 next-hop-self
  exit-address-family
  !
  address-family vpnv6 unicast
  neighbor 3.3.3.3 activate
  neighbor 3.3.3.3 route-reflector-client
  neighbor 3.3.3.3 next-hop-self
  neighbor 5.5.5.5 activate
  neighbor 5.5.5.5 route-reflector-client
  neighbor 5.5.5.5 next-hop-self
  neighbor 8.8.8.8 activate
  neighbor 8.8.8.8 route-reflector-client
  neighbor 8.8.8.8 next-hop-self
  exit-address-family
  !
  exit
```

RR2

BGP

```
bgp extended-asn-cap
!
router bgp 90000
  bgp router-id 7.7.7.7
```

```
no bgp inbound-route-filter
neighbor 3.3.3.3 remote-as 90000
neighbor 3.3.3.3 update-source lo
neighbor 5.5.5.5 remote-as 90000
neighbor 5.5.5.5 update-source lo
neighbor 8.8.8.8 remote-as 90000
neighbor 8.8.8.8 update-source lo
!
address-family vpnv4 unicast
neighbor 3.3.3.3 activate
neighbor 3.3.3.3 route-reflector-client
neighbor 3.3.3.3 next-hop-self
neighbor 5.5.5.5 activate
neighbor 5.5.5.5 route-reflector-client
neighbor 5.5.5.5 next-hop-self
neighbor 8.8.8.8 activate
neighbor 8.8.8.8 route-reflector-client
neighbor 8.8.8.8 next-hop-self
exit-address-family
!
address-family vpnv6 unicast
neighbor 3.3.3.3 activate
neighbor 3.3.3.3 route-reflector-client
neighbor 3.3.3.3 next-hop-self
neighbor 5.5.5.5 activate
neighbor 5.5.5.5 route-reflector-client
neighbor 5.5.5.5 next-hop-self
neighbor 8.8.8.8 activate
neighbor 8.8.8.8 route-reflector-client
neighbor 8.8.8.8 next-hop-self
exit-address-family
!
exit
```

PE2

VRF

```
ip vrf vrf1000
rd 90000:5
route-target both 90000:1
```

VRF Sub interface

```
interface xe2.100
encapsulation dot1q 100
ip vrf forwarding vrf1000
ip address 200.1.1.5/24
ipv6 address 2000::5/64
```

BGP

```
bgp extended-asn-cap
!
router bgp 90000
bgp router-id 5.5.5.5
neighbor 4.4.4.4 remote-as 90000
neighbor 4.4.4.4 update-source lo
neighbor 7.7.7.7 remote-as 90000
neighbor 7.7.7.7 update-source lo
!
address-family vpnv4 unicast
neighbor 4.4.4.4 activate
neighbor 7.7.7.7 activate
```

```
exit-address-family
!
address-family vpnv6 unicast
neighbor 4.4.4.4 activate
neighbor 7.7.7.7 activate
exit-address-family
!
address-family ipv4 vrf vrf1000
redistribute connected
exit-address-family
!
address-family ipv6 vrf vrf1000
redistribute connected
exit-address-family
!
exit
```

PE3

VRF

```
ip vrf vrf1000
rd 90000:3
route-target both 90000:1
```

VRF Sub interface

```
interface xe6.100
encapsulation dot1q 100
ip vrf forwarding vrf1000
ip address 200.1.1.3/24
ipv6 address 2000::3/64
```

BGP

```
bgp extended-asn-cap
!
router bgp 90000
bgp router-id 3.3.3.3
neighbor 4.4.4.4 remote-as 90000
neighbor 4.4.4.4 update-source lo
neighbor 7.7.7.7 remote-as 90000
neighbor 7.7.7.7 update-source lo
!
address-family vpnv4 unicast
neighbor 4.4.4.4 activate
neighbor 7.7.7.7 activate
exit-address-family
!
address-family vpnv6 unicast
neighbor 4.4.4.4 activate
neighbor 7.7.7.7 activate
exit-address-family
!
address-family ipv4 vrf vrf1000
redistribute connected
exit-address-family
!
address-family ipv6 vrf vrf1000
redistribute connected
exit-address-family
!
exit
```


Validation

To verify the ECMP configuration, check the output of the following `show` commands:

```
PE1#show ip bgp vpnv4 all summary
BGP router identifier 8.8.8.8, local AS number 90000
BGP table version is 10
1 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|-------|--------|--------|--------|-----|------|----------|--------------|------|
| 4.4.4.4 | 4 | 90000 | 258 | 243 | 10 | 0 | 0 | 01:28:59 | 2 | |
| 7.7.7.7 | 4 | 90000 | 260 | 249 | 10 | 0 | 0 | 01:28:59 | 2 | |

Total number of neighbors 2

Total number of Established sessions 2

```
PE1#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
Codes: > - installed FTN,
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR
```

| VRF | DESTINATION | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref | Prefix |
|-----|-----------------|-----------------------|---------|--------|------------|--------|
| ID | FEC | cnt | cnt | Ifname | IFNAME | LABEL |
| 2 | B> 200.1.1.0/24 | - | /- | /- | /280 | - |
| | 0x20000801 | 01:29:54 | 2 | 2 | | |
| | | - | /- | /PRI | /266 | - |
| | 26241 | | 4.4.4.4 | | 0x2000cce4 | |
| | - | /- | /PRI | /275 | - | |
| | 26241 | | 7.7.7.7 | | 0x2000cce5 | |
| | - | /- | /PRI | /88 | - | |
| | 26242 | | 4.4.4.4 | | 0x2000cce6 | |
| | - | /- | /PRI | /86 | - | |
| | 26242 | | 7.7.7.7 | | 0x2000cce7 | |
| 2 | B> 2000::/64 | - | /- | /- | /280 | - |
| | 0x20000801 | 01:04:07 | 2 | 2 | | |
| | | - | /- | /PRI | /266 | - |
| | 26241 | | 4.4.4.4 | | 0x2000cce4 | |
| | - | /- | /PRI | /275 | - | |
| | 26241 | | 7.7.7.7 | | 0x2000cce5 | |
| | - | /- | /PRI | /88 | - | |
| | 26242 | | 4.4.4.4 | | 0x2000cce6 | |
| | - | /- | /PRI | /86 | - | |
| | 26242 | | 7.7.7.7 | | 0x2000cce7 | |

```
PE1#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services
```

| Code | FEC | Nexthop | FTN-ID | VRF-ID | Nhlfe-ID | Pri | Out-Label | Out- |
|------|--------------|---------|--------|-----------|----------|-----|-----------|------|
| Intf | | | UpTime | Ext-Color | | | | |
| B> | 200.1.1.0/24 | | 1 | 2 | 280 | (e) | - | - |
| | 01:30:00 | - | | | | | | |
| | 4.4.4.4 | | - | - | 266 | Yes | 26241 | - |
| | 7.7.7.7 | | - | - | 275 | Yes | 26241 | - |
| | 4.4.4.4 | | - | - | 88 | Yes | 26242 | - |

```

      7.7.7.7      -      -      86      Yes      26242      -
B> 2000::/64      3      2      280      (e)      -      -
      01:04:31 -
      4.4.4.4      -      -      266      Yes      26241      -
      7.7.7.7      -      -      275      Yes      26241      -
      4.4.4.4      -      -      88      Yes      26242      -
      7.7.7.7      -      -      86      Yes      26242      -

PE1#sho ip route vrf vrf1000 database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
ia - IS-IS inter area, E - EVPN,
v - vrf leaked
> - selected route, * - FIB route, p - stale info

IP Route Table for VRF "vrf1000"
C  *> 100.1.1.0/24 is directly connected, xe26.100, installed 01:44:08, last update 01:44:08 ago
O  100.1.1.0/24 [110/1] is directly connected, xe26.100, installed 01:43:58, last update
01:43:58 ago
C  *> 127.0.0.0/8 is directly connected, lo.vrf1000, installed 01:44:33, last update 01:44:33 ago
B  > 200.1.1.0/24 [200/0] via 7.7.7.7, installed 01:30:37, last update 01:30:37 ago
    > [200/0] via 7.7.7.7
    > [200/0] via 4.4.4.4
    > [200/0] via 4.4.4.4

Total number of IPv4 routes 7

Gateway of last resort is not set
PE1#

```

Scenarios

Forward Equivalence Class (FEC) is a hardware object to which routes are mapped. It manages the forwarding of packets to an egress port and determines path protection mechanisms.

In the context of FEC, here are a few scenarios:

1. **ECMP FEC resources are full:** This is indicated via oper log and the route/prefix installation is attempted over one single member path. By default, the first member path is selected. If this path is deleted, the next available path is selected.



Note: Logging level must be set correctly.

Log:

```

2025 Jun 19 11:01:51.912 : OcNOS : HSL : NOTIF : [HSL_HWTABLE_FULL_4]:
hardware resource [type : ECMP-FEC] full !!

```

To check the hardware installation status, use the `show hsl mpls l3vpn-ftn detail` command. The path is represented by (e) flag for ECMP and (s) flag for single path.

2. **Per path FEC or MPLS label encapsulation resources are full:** When this occurs, the Next Hop Label Forwarding Entry (NHLFE) object is moved to a retry-table and all the routes/prefixes are skipped from the hardware installation. Once the retry timer which is a default of 10 minutes, expires, the same NHLFE is reattempted in the hardware. To indicate the resource full condition, an oper log is generated.

Logs:

```
2025 Jun 19 11:18:09.082 : OcNOS : HSL : NOTIF : [HSL_HWTABLE_FULL_4]:
hardware resource [type : FEC] full !!
```

```
2025 Jun 19 11:37:23.949 : OcNOS : HSL : NOTIF : [HSL_HWTABLE_FULL_4]:
hardware resource [type : MPLS ENCAP] full !!
```

To check if the prefixes are installed in the hardware, use the `show hsl mpls l3vpn-fts detail` command. It is indicated by > flag.

To check the forwarding entry in the retry table, use the `show hsl mpls nhlf-retry-table` command.

```
OcNOS#show hsl mpls nhlf-retry-table
nhlf-id      flags      error flags      mem nh_list cnt
23           0x30200    0x2              3
```

3. Other scenarios when a similar behavior is seen:

- New member path is added or updated
- Existing member path is deleted



Note: When a tunnel policy configuration is attached to a VRF, the policy is applied to all the routes learned within that VRF. In an ECMP scenario, it is expected that all the ECMP paths are associated with an SR policy as defined by the configured `tunnel-policy`. Due to this, the use of service mapping is not ideal in ECMP. Instead, an SR-ODN or color-based mapping is recommended as a more suitable alternative.

ECMP for L3VPN Commands

The ECMP feature in L3VPN introduces the following configuration commands:

multi-path

Use this command to configure BGP to calculate the multiple routes at RD level for a given address family. Currently, OcNOS supports this command for `vpn4` and `vpn6` unicast address families. When this command is configured, BGP processes the multiple best routes per RD and publishes the routes to multiple VRFs, where the routes get installed based on the BGP config `<max-paths>` at VRF level.

It is recommended to use this command along with `config <max-paths>` at VRF level. This facilitates the ECMP functionality for VPN routes so the best path calculations are done at each RD level and VRF level and the final ECMP paths are chosen at VRF level.

Command Syntax

```
multi-path
no multi-path
```

Parameters

None

Default

Disabled

Command Mode

BGP Address-family L2VPN EVPN

Examples

Consider a topology such that BGP on local PE receives four advertisements from each of its peer PEs, with below IP-VRF parameters:

-RT=200:1, RD=x from PE-1

-RT=200:1, RD=x from PE-2

-RT=200:1, RD=y from PE-3

-RT=200:1, RD=y from PE-4.

Example 1

The following example illustrates the installation of four paths in VRF RED.

Configure a VRF

```
7535-2(config)#ip vrf RED
7535-2(config-vrf)#rd 2:2
7535-2(config-vrf)#route-target both 200:1
7535-2(config-vrf)#exit
```

Configure BGP

```
7535-2(config)#router bgp 100
7535-2(config-router)#address-family l2vpn evpn
7535-2(config-router-af)#multi-path
7535-2(config-router-af)#exit
7535-2(config-router)#address-family ipv4 vrf RED
7535-2(config-router-af)#max-paths ibgp 4
```

Here, BGP learns 2 ECMP paths with RT 200:1 with RD x, and 2 ECMP paths with RT 200:1 with RD y. Finally, VRF RED will have a route with four multipaths installed in the fastpath and this leads to load balancing of user traffic across the four paths.

Example 2

The below example installs two paths in VRF RED. The Multipath configuration is not enabled, however <max-paths> is configured at VRF level.

```
7535-2(config)#router bgp 100
7535-2(config-router)#address-family l2vpn evpn
7535-2(config-router-af)#exit
7535-2(config-router)#address-family ipv4 vrf RED
7535-2(config-router-af)#max-paths ibgp 4
```

Here, BGP learns 2 ECMP paths with RT 200:1 with RD x, and 2 ECMP paths with RT 200:1 with RD y. BGP will not calculate the multipaths per RD, and it will select only the best (one path) per RD. Finally, VRF RED will have a route with two multipaths installed in the fastpath and this leads to load balancing of user traffic across the two paths.

To unconfigure the BGP EVPN Multipath functionality, both <multi-path> and <max-path> configurations must be removed.

load-balance multi-level-ecmp

Use this command to enable multi-level load balancing. For example, L3VPN-ECMP over RSVP-Multipath/LDP ECMP.

Use `no` form of the command to disable the configuration.

Command Syntax

```
load-balance multi-level-ecmp
no load-balance multi-level-ecmp
```

Parameters

None

Default

Disabled

Command Mode

Config Mode

Examples

```
OcNOS#configure terminal
```

Enter the configuration commands ending with `CNTRL/Z`.

```
OcNOS(config)#load-balance multi-level-ecmp
```

Below are the revised `show` commands. For more details, refer to the [MPLS Commands \(page 1381\)](#) chapter.

[show hsl mpls l3vpn-fts \(page 1438\)](#)

[show mpls vrf \(page 1474\)](#)

[show mpls vrf-forwarding-table vrf \(page 1475\)](#)

ECMP Support for L3EVPN

Overview

Equal Cost Multipath (ECMP) for L3EVPN helps in load-balancing the EVPN IRB traffic. It supports traffic on all the IP paths that are available in an MPLS based EVPN network with a symmetric IRB (S-IRB) configuration. This feature programs IP prefix advertisements received on multiple BGP paths into the forwarding plane, enabling ECMP load balancing of known unicast inter-subnet IP traffic.


Feature Characteristics

This feature supports the MPLS-based symmetric IRB (S-IRB) configuration between peer PEs to configure IP multipaths in forwarding-plane.

It applies specifically to MPLS-based EVPN deployments on Qumran1, Qumran2, and J2C+ platforms using VLAN-based or VLAN-aware bundle service interfaces between CE and PE. It ensures that hashing and L3 traffic load-balancing occurs at the ingress PE on a per-service level, independent of the transport-level hashing.

In the below image, the PE routers PE1 and PE2 learn IPv4 and IPv6 customer subnet routes from connected CE device into the corresponding EVPN MAC-VRF via VLAN-based or VLAN-aware bundle service interface. These PE devices' configurations have MAC-VRF connected to an IP-VRF via IRB interface. The protocol BGP running on these PE devices advertise the IRB IP/MAC, CE host IP/MAC, and customer subnet IP routes to UPE as Type-2 and Type-5 (carrying L3VNI for VXLAN-based EVPN or Layer-3 label for MPLS-based EVPN) routes containing the respective IP and MAC-VRF Route Distinguishers and Route Targets.

Figure 95. L3EVPN ECMP Characteristics



The User-facing PE (UPE) processes the advertised routes and their next-hops in its control-plane and installs them in its forwarding-plane as multipath route entries. It prepares the hashing keys based on the configured fields of the incoming traffic, and also load balances the traffic at EVPN service level (either Layer 2 or Layer 3) and independently at transport level (LDP/RSVP/SR-MPLS).

BGP calculates the multipath routes at both RD and VRF levels and passes the route with multiple next-hops towards NSM. NSM and HSL use the existing ECMP infrastructure to install the EVPN routes with multipaths in fast path. HSL uses hierarchical Forwarding Equivalence Class (FEC) and load balances the traffic at service level, within each ECMP group among its multipath member routes.

Benefits

Multipath programming of EVPN Type-2 and Type-5 BGP multipath routes into the data plane.

Operates over MPLS transport with support for LDP, RSVP-TE, or SR-MPLS.

L3/L4 load balancing of inter-subnet known unicast traffic at EVPN IRB service level.

Limitations

Support for symmetric IRB configuration only.

Compatible with P2P transport tunnels only.

Independent L2 and L3 load-balancing—this feature focuses only on L3 (IRB) ECMP.

Supported Hardware

Qumran1, Qumran2, and J2C+ devices.

Prerequisites

1. Define Interfaces and Loopback Addresses

Configure Layer 2 interfaces such as port channel interfaces, for example po1, and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity.

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 8.8.8.8/32 secondary  
  ipv6 address ::1/128  
  
interface po7  
  ip address 31.1.1.8/24
```

2. Configure IGP for Dynamic Routing

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

ISIS Configuration

```
router isis 1  
  is-type level-2-only  
  metric-style wide  
  microloop-avoidance level-2  
  mpls traffic-eng router-id 8.8.8.8  
  mpls traffic-eng level-2  
  capability cspf  
  dynamic-hostname  
  bfd all-interfaces  
  net 49.0000.0000.0008.00  
  passive-interface lo  
!  
interface po7  
  mpls ldp-igp sync isis level-2  
  isis network point-to-point  
  ip router isis 1
```

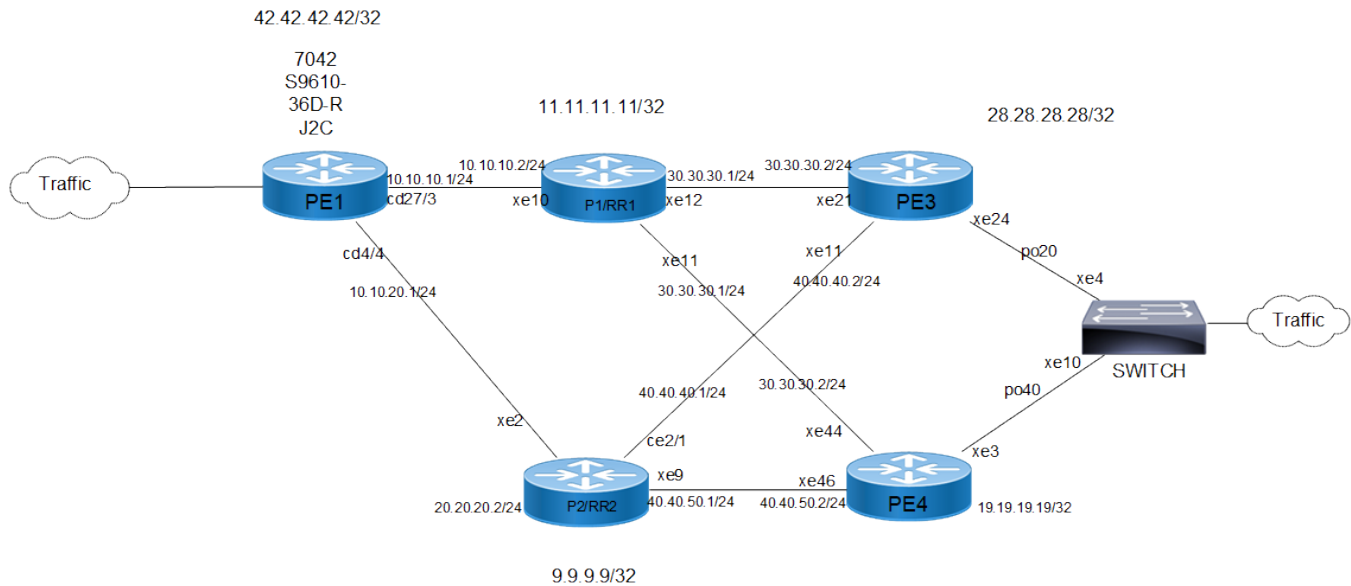
Configuration

To set up ECMP for L3EVPN, follow the configuration steps mentioned below.

Topology

This topology shows the traffic distribution and load balancing across the equal cost paths from PE1 with EVPN service configured in symmetric IRB mode across all the PEs.

The LDP sessions are established between all P and PE routers.

Figure 96. L3EVPN ECMP Topology

PE1

1. Configure Router LDP

```
PE1(config)#router ldp
PE1(config-router)# router-id 8.8.8.8
PE1(config-router)# transport-address ipv4 8.8.8.8
PE1(config-router)#
```

2. Enable LDP and label-switching for core interface

```
PE1(config)#interface cd4
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching

PE1(config)#interface cd27
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching
```

3. Configure VRF

```
PE1(config)#ip vrf ip_vrfirb
PE1(config-vrf)# rd 90000:8
PE1(config-vrf)# route-target both 90000:1
PE1(config-vrf)#
```

4. Configure subinterface and attach VRF to subinterface

```
PE1(config)#interface xe26.100
PE1(config-if)# encapsulation dot1q 100
PE1(config-if)# ip vrf forwarding ip_vrfirb
PE1(config-if)# ip address 100.1.1.1/24
PE1(config-if)# ipv6 address 1000::1/24
```

5. Configure BGP

```
PE1(config)#router bgp 90000
PE1(config-router)# bgp router-id 8.8.8.8
PE1(config-router)# neighbor 4.4.4.4 remote-as 90000
PE1(config-router)# neighbor 4.4.4.4 update-source lo
```



```

PE1(config-router)# neighbor 7.7.7.7 remote-as 90000
PE1(config-router)# neighbor 7.7.7.7 update-source lo
PE1(config-router)# !
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# address-family ipv4 vrf ip_vrfirb
PE1(config-router-af)# redistribute connected
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# address-family ipv6 vrf ip_vrfirb
PE1(config-router-af)# redistribute connected
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# exit
PE1(config)#

```

6. Configure Multipath and max-paths iBGP

```

PE1(config)#router bgp 90000
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#multi-path
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family ipv4 vrf ip_vrfirb
PE1(config-router-af)#max-paths ibgp 4
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family ipv6 vrf ip_vrfirb
PE1(config-router-af)#max-paths ibgp 4
PE1(config-router-af)#
PE1(config-router-af)#exit
PE1(config-router)#

```

P1/RR1

1. Configure Router LDP

```

P1(config)#router ldp
P1(config-router)# router-id 4.4.4.4
P1(config-router)# transport-address ipv4 4.4.4.4
P1(config-router)#

```

2. Enable LDP and label-switching for core interface

```

P1(config)#interface xe12
P1(config-if)# enable-ldp ipv4
P1(config-if)#label-switching

P1(config)#interface xe10
P1(config-if)# enable-ldp ipv4
P1(config-if)#label-switching

P1(config)#interface xe11
P1(config-if)# enable-ldp ipv4
P1(config-if)#label-switching

```

3. Configure BGP

```

P1(config)#router bgp 90000
P1(config-router)# bgp router-id 4.4.4.4
P1(config-router)# no bgp inbound-route-filter
P1(config-router)# neighbor 3.3.3.3 remote-as 90000
P1(config-router)# neighbor 3.3.3.3 update-source lo
P1(config-router)# neighbor 5.5.5.5 remote-as 90000
P1(config-router)# neighbor 5.5.5.5 update-source lo
P1(config-router)# neighbor 8.8.8.8 remote-as 90000
P1(config-router)# neighbor 8.8.8.8 update-source lo

```

```
P1(config-router)# !
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
P1(config-router)#
```

P2

1. Configure Router LDP

```
P2(config)#router ldp
P2(config-router)# router-id 7.7.7.7
P2(config-router)# transport-address ipv4 7.7.7.7
P2(config-router)#
```

2. Enable LDP and label-switching for core interface

```
P2(config)#interface xe2
P2(config-if)# enable-ldp ipv4
P2(config-if)#label-switching
```

```
P2(config)#interface xe9
P2(config-if)# enable-ldp ipv4
P2(config-if)#label-switching
```

```
P2(config)#interface ce2
P2(config-if)# enable-ldp ipv4
P2(config-if)#label-switching
```

3. Configure BGP

```
P2(config)#router bgp 90000
P2(config-router)# bgp router-id 7.7.7.7
P2(config-router)# no bgp inbound-route-filter
P2(config-router)# neighbor 3.3.3.3 remote-as 90000
P2(config-router)# neighbor 3.3.3.3 update-source lo
P2(config-router)# neighbor 5.5.5.5 remote-as 90000
P2(config-router)# neighbor 5.5.5.5 update-source lo
P2(config-router)# neighbor 8.8.8.8 remote-as 90000
P2(config-router)# neighbor 8.8.8.8 update-source lo
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
P2(config-router)#
```

PE3

1. Configure Router LDP

```
PE3(config)#router ldp
PE3(config-router)# router-id 5.5.5.5
PE3(config-router)# transport-address ipv4 5.5.5.5
PE3(config-router)#
```

2. Enable LDP and label-switching for core interface

```
PE3(config)#interface xe21
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching
```

```
PE3(config)#interface xe24
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching
```

3. Configure VRF

```
PE3(config)#ip vrf vrf1000
PE3(config-vrf)# rd 90000:5
PE3(config-vrf)# route-target both 90000:1
PE3(config-vrf)#
```

4. Configure sub-interface and attach the VRF to sub-interface

```
PE3(config)#interface xe2.100
PE3(config-if)# encapsulation dot1q 100
PE3(config-if)# ip vrf forwarding ip_vrfirb
PE3(config-if)# ip address 200.1.1.5/24
PE3(config-if)# ipv6 address 2000::5/64
PE3(config-if)#
```

5. Configure BGP

```
PE3(config)#router bgp 90000
PE3(config-router)# bgp router-id 5.5.5.5
PE3(config-router)# neighbor 4.4.4.4 remote-as 90000
PE3(config-router)# neighbor 4.4.4.4 update-source lo
PE3(config-router)# neighbor 7.7.7.7 remote-as 90000
PE3(config-router)# neighbor 7.7.7.7 update-source lo
PE3(config-router)# !
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE3(config-router)# address-family ipv4 vrf ip_vrfirb
PE3(config-router-af)# redistribute connected
PE3(config-router-af)# exit-address-family
PE3(config-router)# !
PE3(config-router)# address-family ipv6 vrf ip_vrfirb
PE3(config-router-af)# redistribute connected
PE3(config-router-af)# exit-address-family
PE3(config-router)#
```

PE4

1. Configure Router LDP

```
PE4(config)#router ldp
PE4(config-router)# router-id 3.3.3.3
PE4(config-router)# transport-address ipv4 3.3.3.3
PE4(config-router)#
```

2. Enable LDP and label-switching for core interface

```
PE4(config)#interface xe3
PE4(config-if)# enable-ldp ipv4
PE4(config-if)#label-switching
PE4(config)#interface xe46
PE4(config-if)# enable-ldp ipv4
PE4(config-if)#label-switching
```

3. Configure VRF

```
PE4(config)#ip vrf vrf1000
PE4(config-vrf)# rd 90000:3
PE4(config-vrf)# route-target both 90000:1
PE4(config-vrf)#
```

4. Configure sub-interface and attach the VRF to sub-interface

```
PE4(config)#interface xe6.100
PE4(config-if)# encapsulation dot1q 100
PE4(config-if)# ip vrf forwarding ip_vrfirb
PE4(config-if)# ip address 200.1.1.3/24
PE4(config-if)# ipv6 address 2000::3/64
```

```
PE4(config-if)#
```

5. Configure BGP

```
PE4(config)#router bgp 90000
PE4(config-router)# bgp router-id 3.3.3.3
PE4(config-router)# neighbor 4.4.4.4 remote-as 90000
PE4(config-router)# neighbor 4.4.4.4 update-source lo
PE4(config-router)# neighbor 7.7.7.7 remote-as 90000
PE4(config-router)# neighbor 7.7.7.7 update-source lo
PE4(config-router)# !
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE4(config-router)# address-family ipv4 vrf ip_vrfirb
PE4(config-router-af)# redistribute connected
PE4(config-router-af)# exit-address-family
PE4(config-router)# !
PE4(config-router)# address-family ipv6 vrf ip_vrfirb
PE4(config-router-af)# redistribute connected
PE4(config-router-af)# exit-address-family
PE4(config-router)#
```

Configuration Snapshot

PE1

```
PE1-7042#sh running-config
!
! Software version: UFI_S9610-36D-OcNOS-SP-PLUS-7.0.0.136-Alpha 10/12/2025 17:35:02
!
! Last configuration change at 18:47:41 UTC Mon Oct 13 2025 by root
!
!
service password-encryption
!
no logging cli
logging logfile rsvp 7
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE1-7042
port cd4 breakout 4X10g
port cd6 breakout 4X10g
port cd27 breakout 4X10g
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
```

```
lldp notification-interval 1000
fault-management enable
!
evpn mpls enable
!
evpn mpls irb
!
ip vrf ip_vrfirb
  rd 42.42.42.42:200
  route-target both 200:200
  l3vni 20000
!
ip vrf ip_vrfirbMH801
  rd 42.42.42.42:801
  route-target both 801:801
  l3vni 40000
!
ip vrf management
!
mac vrf vrfirb
  rd 42.42.42.42:2000
  route-target both 2000:2000
!
mac vrf vrfirbMH801
  rd 42.42.42.42:4000
  route-target both 4000:4000
!
evpn mpls vtep-ip-global 42.42.42.42
!
evpn mpls id 200
  host-reachability-protocol evpn-bgp vrfirb
!
evpn mpls id 801
  host-reachability-protocol evpn-bgp vrfirbMH801
  evpn irb irb700
!
router rsvp
  auto-bypass
  attributes best-effort
  protection-capability node
  reoptimize
  exit
  inactivity-timer 5
  enable
  exit
!
interface cd0
!
interface cd1
!
interface cd2
!
interface cd3
!
interface cd4/1
!
interface cd4/2
!
interface cd4/3
!
interface cd4/4
  load-interval 30
  ip address 10.10.20.1/24
  mtu 9216
  label-switching
  shutdown
  ip router isis 1
  enable-rsvp
```

```
!  
interface cd5  
!  
interface cd6/1  
!  
interface cd6/2  
!  
interface cd6/3  
!  
interface cd6/4  
!  
interface cd6/4.200 switchport  
    encapsulation dot1q 200  
    rewrite pop  
    access-if-evpn  
    map vpn-id 200  
!  
interface cd6/4.801 switchport  
    encapsulation dot1q 801  
    rewrite pop  
    access-if-evpn  
    map vpn-id 801  
!  
load-balance enable  
!  
load-balance ipv4 src-dest-ipv4  
    load-balance ipv6 flow-label next-hdr src-dest-ipv6  
    load-balance src-dest-l4port  
!  
interface cd7  
!  
interface cd8  
!  
interface cd9  
!  
interface cd10  
!  
interface cd11  
!  
interface cd12  
!  
interface cd13  
!  
interface cd14  
!  
interface cd15  
!  
interface cd16  
!  
interface cd17  
!  
interface cd18  
!  
interface cd19  
!  
interface cd20  
!  
interface cd21  
!  
interface cd22  
!  
interface cd23  
!  
interface cd24  
!  
interface cd25  
!  
interface cd26
```

```
!  
interface cd27/1  
!  
interface cd27/2  
!  
interface cd27/3  
  load-interval 30  
  ip address 10.10.10.1/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-rsvp  
!  
interface cd27/4  
!  
interface cd28  
!  
interface cd29  
!  
interface cd30  
!  
interface cd31  
!  
interface cd32  
!  
interface cd33  
!  
interface cd34  
!  
interface cd35  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface irb100  
  ip vrf forwarding ip_vrfirb  
  ip address 80.80.1.1/24 anycast  
  ipv6 address 80::1/64 anycast  
!  
interface irb700  
  ip vrf forwarding ip_vrfirbMH801  
  ip address 81.81.1.1/24 anycast  
  ipv6 address 81::1/64 anycast  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 42.42.42.42/32 secondary  
  ipv6 address ::1/128  
  ip router isis 1  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
evpn irb-forwarding anycast-gateway-mac 0011.2200.0001  
!  
  exit  
!  
router ospf  
!  
router ospf 100  
  bfd all-interfaces  
  network 10.10.10.0/24 area 0.0.0.0  
  network 10.10.20.0/24 area 0.0.0.0  
  network 42.42.42.42/32 area 0.0.0.0  
!
```

```

router bgp 65100
  bgp router-id 42.42.42.42
  no bgp inbound-route-filter
  allocate-label all
  neighbor 11.11.11.11 remote-as 65100
  neighbor 11.11.11.11 update-source lo
  neighbor 81.81.1.2 remote-as 65101
  !
  address-family ipv4 unicast
  network 42.42.42.42/32
  exit-address-family
  !
  address-family l2vpn evpn
  multi-path
  neighbor 11.11.11.11 activate
  neighbor 11.11.11.11 route-reflector-client
  exit-address-family
  !
  address-family ipv4 vrf ip_vrfirb
  redistribute connected
  exit-address-family
  !
  address-family ipv4 vrf ip_vrfirbMH801
  max-paths ibgp 4
  redistribute connected
  neighbor 81.81.1.2 remote-as 65101
  neighbor 81.81.1.2 activate
  exit-address-family
  !
  exit
  !
  rsvp-path PE1-PE3 mpls
  10.10.10.2 strict
  30.30.40.2 strict
  !
  rsvp-path PE1-PE4 mpls
  10.10.10.2 strict
  30.30.30.2 strict
  !
  rsvp-trunk PE1-PE2 ipv4
  to 30.30.30.30
  !
  rsvp-trunk PE1-PE3 ipv4
  no primary cspf
  primary path PE1-PE3
  to 28.28.28.28
  !
  rsvp-trunk PE1-PE4 ipv4
  no primary cspf
  primary path PE1-PE4
  to 19.19.19.19
  !
  !
end

!
PE1-7042#

```

```
PE1-7042#sh ip ospf neighbor
```

```
Total number of full neighbors: 0
OSPF process 0 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------|-----------|---------|-----------|-------------|
|-------------|-----|-------|-----------|---------|-----------|-------------|

```
Total number of full neighbors: 1
```


OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 11.11.11.11 | 1 | Full/Backup | 00:00:34 | 10.10.10.2 | cd27/3 | 0 |

PE1-7042#

PE1-7042#sh rsvp session

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass

State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary

* indicates the session is active with local repair at one or more nodes

(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|---------------|----------|--------|------|----------|-------|
| Uptime Rt | Style Labelin | Labelout | | | | |
| 30.30.30.30 | 42.42.42.42 | 5001 | 2201 | PRI | PE1-PE2- | |
| Primary | UP | 00:20:33 | 1 1 SE | - | 24327 | |
| 28.28.28.28 | 42.42.42.42 | 5002 | 2203 | PRI | PE1-PE3- | |
| Primary | UP | 00:21:28 | 1 1 SE | - | 24326 | |
| 19.19.19.19 | 42.42.42.42 | 5003 | 2204 | PRI | PE1-PE4- | |
| Primary | UP | 00:21:28 | 1 1 SE | - | 24325 | |

Total 3 displayed, Up 3, Down 0.

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|---------------|----------|--------|-------|----------|-------|
| Uptime Rt | Style Labelin | Labelout | | | | |
| 42.42.42.42 | 19.19.19.19 | 5001 | 2201 | PRI | PE4-PE1- | |
| Primary | UP | 00:20:26 | 1 1 SE | 24960 | - | |
| 42.42.42.42 | 28.28.28.28 | 5001 | 2201 | PRI | PE3-PE1- | |
| Primary | UP | 00:20:01 | 1 1 SE | 24962 | - | |
| 42.42.42.42 | 30.30.30.30 | 5001 | 2202 | PRI | PE2-PE1- | |
| Primary | UP | 00:20:21 | 1 1 SE | 24961 | - | |

Total 3 displayed, Up 3, Down 0.

PE1-7042#

PE1-7042#

PE1-7042#show bgp l2vpn evpn

BGP table version is 4, local router ID is 42.42.42.42

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,

l - labeled, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]

1 - Ethernet Auto-discovery Route

2 - MAC/IP Route

3 - Inclusive Multicast Route

4 - Ethernet Segment Route

5 - Prefix Route

| Network | Next Hop | Metric | LocPrf | Weight | Path | Peer | Encap |
|--------------------------------|-------------|--------|--------|--------|------|---------------|-------|
| RD[65100:501] | | | | | | | |
| *>i [3]:[501]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 | i | - 11.11.11.11 | MPLS |
| RD[65100:502] | | | | | | | |
| *>i [3]:[502]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 | i | - 11.11.11.11 | MPLS |
| RD[65100:503] | | | | | | | |
| *>i [3]:[503]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 | i | - 11.11.11.11 | MPLS |
| RD[65100:504] | | | | | | | |
| *>i [3]:[504]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 | i | - 11.11.11.11 | MPLS |
| RD[65100:505] | | | | | | | |
| *>i [3]:[505]:[32,30.30.30.30] | | | | | | | |

```

30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS

RD[65100:506]
*>i      [3]:[506]:[32,30.30.30.30]
30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS

RD[65100:507]
*>i      [3]:[507]:[32,30.30.30.30]
30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS

RD[65100:508]
*>i      [3]:[508]:[32,30.30.30.30]
30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS

RD[65100:509]
*>i      [3]:[509]:[32,30.30.30.30]
30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS

RD[65100:510]
*>i      [3]:[510]:[32,30.30.30.30]
30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS

RD[19.19.19.19:250]
*>i      [5]:[0]:[0]:[24]:[70.70.1.0]:[0.0.0.0]:[16]
19.19.19.19      0      100      0      ?      -  11.11.11.11      MPLS

RD[19.19.19.19:801]
*>i      [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
19.19.19.19      0      100      0      ?      -  11.11.11.11      MPLS

RD[19.19.19.19:3000]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[251]:[19]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1]:[19]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS
*>i      [3]:[251]:[32,19.19.19.19]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS

RD[19.19.19.19:4000]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS
*>i      [3]:[811]:[32,19.19.19.19]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS

RD[19.19.19.19:64512]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS
*>i      [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
19.19.19.19      0      100      0      i      -  11.11.11.11      MPLS

RD[28.28.28.28:200]
*>i      [5]:[0]:[0]:[24]:[90.90.1.0]:[0.0.0.0]:[16]
28.28.28.28      0      100      0      ?      -  11.11.11.11      MPLS

RD[28.28.28.28:801]
*>i      [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
28.28.28.28      0      100      0      ?      -  11.11.11.11      MPLS

RD[28.28.28.28:2000]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]

```

```

28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i  [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1][19]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i  [3]:[201]:[32,28.28.28.28]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS

RD[28.28.28.28:4000]
*>i  [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i  [2]:[00:00:00:11:11:22:22:00:00:00]:[811]:[48,0010:9400:0005]:[0]:[20]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i  [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i  [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1][20]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i  [3]:[811]:[32,28.28.28.28]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS

RD[28.28.28.28:64512]
*>i  [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i  [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS

RD[30.30.30.30:250]
*>i  [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
30.30.30.30      0      100      0      ?      -      11.11.11.11      MPLS

RD[30.30.30.30:3000]
*>i  [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
30.30.30.30      0      100      0      i      -      11.11.11.11      MPLS
*>i  [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1][19]
30.30.30.30      0      100      0      i      -      11.11.11.11      MPLS
*>i  [3]:[250]:[32,30.30.30.30]
30.30.30.30      0      100      0      i      -      11.11.11.11      MPLS

RD[42.42.42.42:2000] VRF[vrfirb]
* i  [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
* i  [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>  [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
42.42.42.42      0      100      32768  i      -      -----
MPLS
*>  [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
42.42.42.42      0      100      32768  i      -      -----
MPLS
*>  [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]
42.42.42.42      0      100      32768  i      -      -----
MPLS
*>  [2]:[0]:[200]:[48,1200:0100:0003]:[0]:[18]
42.42.42.42      0      100      32768  i      -      -----
MPLS
*>  [2]:[0]:[200]:[48,1200:0100:0004]:[0]:[18]
42.42.42.42      0      100      32768  i      -      -----
MPLS
*>  [2]:[0]:[200]:[48,1200:0100:0005]:[0]:[18]
42.42.42.42      0      100      32768  i      -      -----
MPLS
*>  [2]:[0]:[200]:[48,1200:0100:0006]:[0]:[18]
42.42.42.42      0      100      32768  i      -      -----
MPLS
*>  [2]:[0]:[200]:[48,1200:0100:0007]:[0]:[18]
42.42.42.42      0      100      32768  i      -      -----
MPLS
*>  [2]:[0]:[200]:[48,1200:0100:0008]:[0]:[18]
42.42.42.42      0      100      32768  i      -      -----
MPLS
*>  [2]:[0]:[200]:[48,1200:0100:0009]:[0]:[18]

```

```

42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000a]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000b]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000c]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000d]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000e]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000f]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0010]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0011]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0012]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0013]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0014]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0015]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0016]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0017]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0018]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0019]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001a]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001b]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001c]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001d]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001e]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001f]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0020]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS

```

```

*> [2]:[0]:[200]:[48,1200:0100:0021]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0022]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0023]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0024]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0025]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0026]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0027]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0028]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0029]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002a]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002b]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002c]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002d]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002e]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002f]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0030]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0031]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0032]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0033]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0034]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0035]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0036]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0037]:[0]:[18]
      42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0038]:[0]:[18]

```

```

42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0039]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003a]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003b]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003c]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003d]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003e]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003f]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0040]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0041]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0042]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0043]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0044]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0045]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0046]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0047]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0048]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0049]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004a]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004b]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004c]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004d]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004e]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004f]:[0]:[18]
42.42.42.42      0      100      32768 i      -      -----
MPLS

```

```

*> [2]:[0]:[200]:[48,1200:0100:0050]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0051]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0052]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0053]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0054]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0055]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0056]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0057]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0058]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0059]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
* i [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
* i [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1]:[19]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
*> [3]:[200]:[32,42.42.42.42]
      42.42.42.42      0      100      32768 i      - -----
MPLS

```

```

* i   [3]:[201]:[32,28.28.28.28]
      28.28.28.28          0          100          0   i   - 11.11.11.11      MPLS

RD[42.42.42.42:4000] VRF[vrfirbMH801]
* i   [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28          0          100          0   i   - 11.11.11.11      MPLS
* i   19.19.19.19          0          100          0   i   - 11.11.11.11      MPLS
* i   [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28          0          100          0   i   - 11.11.11.11      MPLS
* i   19.19.19.19          0          100          0   i   - 11.11.11.11      MPLS
*>   [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
      42.42.42.42          0          100          32768 i   - -----
      MPLS
*>   [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1]:[19]
      42.42.42.42          0          100          32768 i   - -----
      MPLS
*>   [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
      42.42.42.42          0          100          32768 i   - -----
      MPLS
*>   [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
      42.42.42.42          0          100          32768 i   - -----
      MPLS
* i   [2]:[00:00:00:11:11:22:22:00:00:00]:[811]:[48,0010:9400:0005]:[0]:[20]
      28.28.28.28          0          100          0   i   - 11.11.11.11      MPLS
* i   [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28          0          100          0   i   - 11.11.11.11      MPLS
* i   19.19.19.19          0          100          0   i   - 11.11.11.11      MPLS
* i   [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
      28.28.28.28          0          100          0   i   - 11.11.11.11      MPLS
* i   19.19.19.19          0          100          0   i   - 11.11.11.11      MPLS
*>   [3]:[801]:[32,42.42.42.42]
      42.42.42.42          0          100          32768 i   - -----
      MPLS
* i   [3]:[811]:[32,19.19.19.19]
      19.19.19.19          0          100          0   i   - 11.11.11.11      MPLS
* i   [3]:[811]:[32,28.28.28.28]
      28.28.28.28          0          100          0   i   - 11.11.11.11      MPLS

Total number of prefixes 158
PE1-7042#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code   FEC
Label  Out-Intf      Nexthop      UpTime      FTN-ID VRF-ID  Nhlfe-ID  Pri  Out-
      Ext-Color
B(x)>   90.90.1.0/24
      -          00:20:35 -          1          2          13          -  -
      28.28.28.28 -          -          12          Yes  16
B(x)>   91.91.1.0/24
      -          00:20:35 -          1          3          15          (e) -
      19.19.19.19 -          -          10          Yes  17
      28.28.28.28 -          -          14          Yes  17

PE1-7042#show ip route vrf all database bgp
IP Route Table for VRF "default"

Total number of IPv4 routes 0
IP Route Table for VRF "management"

Total number of IPv4 routes 0
IP Route Table for VRF "ip_vrfirb"
B      > 90.90.1.0/24 [200/0] via 28.28.28.28, installed 00:20:35, last update 00:20:35 ago

Total number of IPv4 routes 1

```



```
IP Route Table for VRF "ip_vrfirbMH801"
B    > 91.91.1.0/24 [200/0] via 28.28.28.28, installed 00:20:35, last update 00:20:35 ago
      >                                     [200/0] via 19.19.19.19

Total number of IPv4 routes 2
IP Route Table for VRF "vrfirb"

Total number of IPv4 routes 0
IP Route Table for VRF "vrfirbMH801"

Total number of IPv4 routes 0

Gateway of last resort is not set
PE1-7042#show hsl mpls l3vpn-fts
TABLE: L3VPN Ftn table
Codes: > - installed Ftn, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF |      DESTINATION      | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref |Prefix|
| MPLS |      OUT      | OUT      | NEXTHOP      |         |     |      |
| ID  |      FEC      | cnt | cnt | Ifname      | IFNAME | LABEL |
| OBJECT | cnt | cnt |         |         |         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
2    B> 90.90.1.0/24      13(s) - - - -
      0x2000ccdb 00:22:34 1 1
      - /- /PRI /12 -
      16      28.28.28.28 0x2000ccdb
3    B> 91.91.1.0/24      15(e) - - - -
      0x20000001 00:22:34 1 1
      - /- /PRI /10 -
      17      19.19.19.19 0x2000ccdc
      - /- /PRI /14 -
      17      28.28.28.28 0x2000ccdd

PE1-7042#sh evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      evpn-id      Local-
Leaf Remote-Leaf Ext-Color FAT
=====
42.42.42.42      19.19.19.19      Installed      00:22:37      00:22:37      40000      ---
---      ---      ---
42.42.42.42      28.28.28.28      Installed      00:22:37      00:22:37      40000      ---
---      ---      ---
42.42.42.42      28.28.28.28      Installed      00:22:37      00:22:37      20000      ---
---      ---      ---

Total number of entries are 3
PE1-7042#

apps-fileview.texmex_20251030.00_p0
PE1-7042.txt
Displaying PE1-7042.txt.
```

P1/RR1

```
P1-7011#sh running-config
!
! Software version: HFCL-AR-OcNOS-CSR-7.0.0.133-Alpha
10/09/20
!
! Last configuration change at 11:38:04 UTC Sat Aug 02 2025
by
!
!
service password-encryption
```

```
!  
logging console 5  
logging monitor disable  
logging logfile rsvp 7  
logging level nsm 4  
logging level rsvp 7  
logging level bgp 7  
logging level cmm 4  
logging level all 5  
!  
!  
snmp-server enable traps link linkDown  
snmp-server enable traps link linkUp  
!  
hardware-profile statistics ingress-acl enable  
!  
bfd interval 3 minrx 3 multiplier 3  
!  
qos enable  
!  
hostname P1-7011  
ip name-server vrf management 10.16.10.23  
tfo Disable  
errdisable cause stp-bpdu-guard  
feature dns relay  
ip dns relay  
ipv6 dns relay  
lldp run  
lldp tlv-select basic-mgmt port-description  
lldp tlv-select basic-mgmt system-name  
lldp tlv-select basic-mgmt system-capabilities  
lldp tlv-select basic-mgmt system-description  
lldp tlv-select basic-mgmt management-address  
lldp notification-interval 1000  
fault-management enable  
!  
ip vrf management  
!  
router ldp  
!  
router rsvp  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface ge1  
!  
interface ge2  
!  
interface ge3  
!  
interface ge4  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 11.11.11.11/32 secondary  
  ipv6 address ::1/128  
  ip router isis 1  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xe5  
!  
interface xe6
```

```
!  
interface xe7  
  load-interval 30  
  ip address 20.20.30.2/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface xe8  
!  
interface xe9  
!  
interface xe10  
  load-interval 30  
  ip address 10.10.10.2/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface xe11  
  load-interval 30  
  ip address 30.30.30.1/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface xe12  
  load-interval 30  
  ip address 30.30.40.1/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-ldp ipv4  
  enable-rsvp  
!  
  exit  
!  
router ospf 100  
  bfd all-interfaces  
  network 10.10.10.0/24 area 0.0.0.0  
  network 11.11.11.11/32 area 0.0.0.0  
  network 20.20.30.0/24 area 0.0.0.0  
  network 30.30.30.0/24 area 0.0.0.0  
  network 30.30.40.0/24 area 0.0.0.0  
!  
router isis 1  
  is-type level-2-only  
  metric-style wide  
  microloop-avoidance level-2  
  mpls traffic-eng router-id 11.11.11.11  
  mpls traffic-eng level-2  
  dynamic-hostname  
  fast-reroute per-prefix level-2 proto ipv4 all  
  fast-reroute ti-lfa level-2 proto ipv4  
  bfd all-interfaces  
  net 49.0000.0000.0011.00  
!  
router bgp 65100  
  bgp router-id 11.11.11.11  
  no bgp inbound-route-filter  
  allocate-label all  
  neighbor 19.19.19.19 remote-as 65100
```

```

neighbor 19.19.19.19 update-source lo
neighbor 28.28.28.28 remote-as 65100
neighbor 28.28.28.28 update-source lo
neighbor 30.30.30.30 remote-as 65100
neighbor 30.30.30.30 update-source lo
neighbor 42.42.42.42 remote-as 65100
neighbor 42.42.42.42 update-source lo
!
address-family l2vpn evpn
neighbor 19.19.19.19 activate
neighbor 19.19.19.19 route-reflector-client
neighbor 28.28.28.28 activate
neighbor 28.28.28.28 route-reflector-client
neighbor 30.30.30.30 activate
neighbor 30.30.30.30 route-reflector-client
neighbor 42.42.42.42 activate
neighbor 42.42.42.42 route-reflector-client
exit-address-family
!
exit
!
!
end

!
P1-7011#

```

```
P1-7011#sh ip ospf neighbor
```

```
Total number of full neighbors: 4
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|---------|-----------|------------|-----------|-------------|
| 42.42.42.42 | 1 | Full/DR | 00:00:35 | 10.10.10.1 | xe10 | 0 |
| 30.30.30.30 | 1 | Full/DR | 00:00:35 | 20.20.30.1 | xe7 | 0 |
| 19.19.19.19 | 1 | Full/DR | 00:00:38 | 30.30.30.2 | xe11 | 0 |
| 28.28.28.28 | 1 | Full/DR | 00:00:39 | 30.30.40.2 | xe12 | 0 |

```
P1-7011#
```

```
P1-7011#sh rsvp session
```

```
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
```

```
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
```

```
* indicates the session is active with local repair at one or more nodes
```

```
(P) indicates the secondary-priority session is acting as primary
```

```
Transit RSVP:
```

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|---------------|----------|--------|-------|----------|-------|
| Uptime Rt | Style Labelin | Labelout | | | | |
| 30.30.30.30 | 42.42.42.42 | 5001 | 2201 | PRI | PE1-PE2- | |
| Primary | UP | 00:25:19 | 1 1 SE | 24327 | 24962 | |
| 42.42.42.42 | 19.19.19.19 | 5001 | 2201 | PRI | PE4-PE1- | |
| Primary | UP | 00:25:12 | 1 1 SE | 24328 | 24960 | |
| 42.42.42.42 | 28.28.28.28 | 5001 | 2201 | PRI | PE3-PE1- | |
| Primary | UP | 00:24:47 | 1 1 SE | 24330 | 24962 | |
| 42.42.42.42 | 30.30.30.30 | 5001 | 2202 | PRI | PE2-PE1- | |
| Primary | UP | 00:25:07 | 1 1 SE | 24329 | 24961 | |
| 28.28.28.28 | 30.30.30.30 | 5002 | 2201 | PRI | PE2-PE3- | |
| Primary | UP | 00:29:36 | 1 1 SE | 24321 | 24960 | |
| 28.28.28.28 | 42.42.42.42 | 5002 | 2203 | PRI | PE1-PE3- | |
| Primary | UP | 00:26:14 | 1 1 SE | 24326 | 24962 | |
| 30.30.30.30 | 19.19.19.19 | 5002 | 2202 | PRI | PE4-PE2- | |
| Primary | UP | 00:27:01 | 1 1 SE | 24323 | 24961 | |
| 30.30.30.30 | 28.28.28.28 | 5002 | 2202 | PRI | PE3-PE2- | |
| Primary | UP | 00:29:42 | 1 1 SE | 24320 | 24960 | |
| 19.19.19.19 | 28.28.28.28 | 5003 | 2203 | PRI | PE3-PE4- | |
| Primary | UP | 00:26:46 | 1 1 SE | 24324 | 24960 | |
| 19.19.19.19 | 42.42.42.42 | 5003 | 2204 | PRI | PE1-PE4- | |
| Primary | UP | 00:26:14 | 1 1 SE | 24325 | 24961 | |
| 28.28.28.28 | 19.19.19.19 | 5003 | 2203 | PRI | PE4-PE3- | |
| Primary | UP | 00:27:10 | 1 1 SE | 24322 | 24961 | |

Total 11 displayed, Up 11, Down 0.

Pl-7011#

Pl-7011#

Pl-7011#show bgp l2vpn evpn

BGP table version is 10, local router ID is 11.11.11.11

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,

l - labeled, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevant route information]

1 - Ethernet Auto-discovery Route

2 - MAC/IP Route

3 - Inclusive Multicast Route

4 - Ethernet Segment Route

5 - Prefix Route

| Network | Next Hop | Metric | LocPrf | Weight | Path | Peer | Encap |
|---|-------------|--------|--------|--------|------|-------------|-------|
| RD[65100:501] | | | | | | | |
| *>i [3]:[501]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[65100:502] | | | | | | | |
| *>i [3]:[502]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[65100:503] | | | | | | | |
| *>i [3]:[503]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[65100:504] | | | | | | | |
| *>i [3]:[504]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[65100:505] | | | | | | | |
| *>i [3]:[505]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[65100:506] | | | | | | | |
| *>i [3]:[506]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[65100:507] | | | | | | | |
| *>i [3]:[507]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[65100:508] | | | | | | | |
| *>i [3]:[508]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[65100:509] | | | | | | | |
| *>i [3]:[509]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[65100:510] | | | | | | | |
| *>i [3]:[510]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 30.30.30.30 | MPLS |
| RD[19.19.19.19:250] | | | | | | | |
| *>i [5]:[0]:[0]:[24]:[70.70.1.0]:[0.0.0.0]:[16] | 19.19.19.19 | 0 | 100 | 0 ? | - | 19.19.19.19 | MPLS |
| RD[19.19.19.19:801] | | | | | | | |
| *>i [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17] | 19.19.19.19 | 0 | 100 | 0 ? | - | 19.19.19.19 | MPLS |

```

RD[19.19.19.19:3000]
*>i  [1]:[00:00:00:11:11:22:22:00:00:00]:[251]:[19]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS
*>i  [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS
*>i  [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1]:[19]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS
*>i  [3]:[251]:[32,19.19.19.19]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS

RD[19.19.19.19:4000]
*>i  [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS
*>i  [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS
*>i  [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS
*>i  [3]:[811]:[32,19.19.19.19]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS

RD[19.19.19.19:64512]
*>i  [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS
*>i  [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
      19.19.19.19      0      100      0      i      -      19.19.19.19      MPLS

RD[28.28.28.28:200]
*>i  [5]:[0]:[0]:[24]:[90.90.1.0]:[0.0.0.0]:[16]
      28.28.28.28      0      100      0      ?      -      28.28.28.28      MPLS

RD[28.28.28.28:801]
*>i  [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
      28.28.28.28      0      100      0      ?      -      28.28.28.28      MPLS

RD[28.28.28.28:2000]
*>i  [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS
*>i  [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS
*>i  [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1]:[19]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS
*>i  [3]:[201]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS

RD[28.28.28.28:4000]
*>i  [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS
*>i  [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS
*>i  [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS
*>i  [3]:[811]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS

RD[28.28.28.28:64512]
*>i  [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS
*>i  [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      -      28.28.28.28      MPLS

RD[30.30.30.30:250]
*>i  [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
      30.30.30.30      0      100      0      ?      -      30.30.30.30      MPLS

RD[30.30.30.30:3000]
*>i  [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
      30.30.30.30      0      100      0      i      -      30.30.30.30      MPLS

```

```

*>i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1][19]
          30.30.30.30          0          100          0    i          - 30.30.30.30    MPLS
*>i    [3]:[250]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          - 30.30.30.30    MPLS

RD[42.42.42.42:801]
*>i    [5]:[0]:[0]:[24]:[81.81.1.0]:[0.0.0.0]:[17]
          42.42.42.42          0          100          0    ?          - 42.42.42.42    MPLS

RD[42.42.42.42:2000]
*>i    [3]:[200]:[32,42.42.42.42]
          42.42.42.42          0          100          0    i          - 42.42.42.42    MPLS

RD[42.42.42.42:4000]
*>i    [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
          42.42.42.42          0          100          0    i          - 42.42.42.42    MPLS
*>i    [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1][19]
          42.42.42.42          0          100          0    i          - 42.42.42.42    MPLS
*>i    [3]:[801]:[32,42.42.42.42]
          42.42.42.42          0          100          0    i          - 42.42.42.42    MPLS

Total number of prefixes 43
Pl-7011#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code    FEC
Label    Out-Intf          Nexthop          UpTime          FTN-ID VRF-ID      Nhlfe-ID      Pri    Out-
Ext-Color
Pl-7011#show ip route vrf all database bgp
IP Route Table for VRF "default"

Total number of IPv4 routes 0
IP Route Table for VRF "management"

Total number of IPv4 routes 0

Gateway of last resort is not set
Pl-7011#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF |      DESTINATION      | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref | Prefix |
| MPLS |      OUT      | OUT      | NEXTHOP      |         |     |         |
| ID |      FEC      |         | Ifname      | IFNAME | LABEL |         |
| OBJECT | cnt | cnt |         |         |         |         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Pl-7011#sh evpn mpls tunnel
% Evpn mpls is not Enabled !
Pl-7011#

apps-fileview.texmex_20251030.00_p0
PlRR.txt
Displaying PlRR.txt.

```

P2/RR2

```

P2-7009#sh running-config
!
! Software version: HFCL-CUAR-OcNOS-SP-PLUS-7.0.0.133-Alpha
10
/09/2025 17:37:22

```

```
!  
! Last configuration change at 14:33:47 UTC Sun Sep 07 2025  
by root  
!  
!  
service password-encryption  
!  
logging console 5  
logging logfile rsvp 7  
logging level nsm 4  
logging level rsvp 7  
logging level bgp 7  
logging level cmm 4  
logging level all 5  
!  
!  
snmp-server enable traps link linkDown  
snmp-server enable traps link linkUp  
!  
hardware-profile statistics voq-full-color enable  
hardware-profile statistics cfm-ccm enable  
!  
qos enable  
!  
hostname P2-7009  
port ce2 breakout 4X10g  
ip name-server vrf management 10.16.10.23  
tfo Disable  
errdisable cause stp-bpdu-guard  
feature dns relay  
ip dns relay  
ipv6 dns relay  
feature rsyslog  
logging remote server 10.16.58.70 5 port 1514 vrf management  
lldp run  
lldp tlv-select basic-mgmt port-description  
lldp tlv-select basic-mgmt system-name  
lldp tlv-select basic-mgmt system-capabilities  
lldp tlv-select basic-mgmt system-description  
lldp tlv-select basic-mgmt management-address  
lldp notification-interval 1000  
fault-management enable  
!  
ip vrf management  
!  
router ldp  
!  
router rsvp  
!  
interface ce1  
!  
interface ce2/1  
load-interval 30  
ip address 40.40.40.1/24  
mtu 9216  
label-switching  
ip router isis 1  
enable-ldp ipv4  
enable-rsvp  
!  
interface ce2/2  
!  
interface ce2/3  
!  
interface ce2/4  
!  
interface ce3  
!
```



```
interface ce4
!
interface ce5
!
interface ce6
!
interface ce7
!
interface ce8
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 9.9.9.9/32 secondary
 ipv6 address ::1/128
 ip router isis 1
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xel
!
interface xe2
 speed 10g
 load-interval 30
 ip address 10.10.20.2/24
 mtu 9216
 label-switching
 ip router isis 1
 enable-ldp ipv4
 enable-rsvp
!
interface xe3
!
interface xe4
!
interface xe5
 load-interval 30
 ip address 20.20.20.2/24
 mtu 9216
 label-switching
 ip router isis 1
 enable-ldp ipv4
 enable-rsvp
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
 speed 10g
 load-interval 30
 ip address 40.40.50.1/24
 mtu 9216
 label-switching
 ip router isis 1
 enable-ldp ipv4
 enable-rsvp
!
interface xe10
!
```

```

interface xe11
!
interface xe12
!
interface xe13
!
interface xe14
!
interface xe15
!
interface xe16
!
interface xe17
!
interface xe18
!
interface xe19
!
interface xe20
!
exit
!
router ospf 100
  bfd all-interfaces
  network 9.9.9.9/32 area 0.0.0.0
  network 10.10.20.0/24 area 0.0.0.0
  network 20.20.20.0/24 area 0.0.0.0
  network 40.40.40.0/24 area 0.0.0.0
  network 40.40.50.0/24 area 0.0.0.0
!
router isis 1
  is-type level-2-only
  metric-style wide
  microloop-avoidance level-2
  mpls traffic-eng router-id 9.9.9.9
  mpls traffic-eng level-2
  dynamic-hostname
  fast-reroute per-prefix level-2 proto ipv4 all
  fast-reroute ti-lfa level-2 proto ipv4
  bfd all-interfaces
  net 49.0000.0000.0009.00
!
!
end

!
P2-7009#

P2-7009#sh ip ospf neighbor

Total number of full neighbors: 3
OSPF process 100 VRF(default):
Neighbor ID      Pri   State           Dead Time   Address        Interface      Instance ID
30.30.30.30      1     Full/DR          00:00:32    20.20.20.1     xe5            0
28.28.28.28      1     Full/DR          00:00:35    40.40.40.2     ce2/1          0
19.19.19.19      1     Full/DR          00:00:32    40.40.50.2     xe9            0
P2-7009#
P2-7009#sh rsvp session
Type   : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State  : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

P2-7009#
P2-7009#
P2-7009#show bgp l2vpn evpn

```

```

Total number of prefixes 0
P2-7009#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code      FEC
Label      Out-Intf      Nexthop      UpTime      FTN-ID VRF-ID      Nhlfe-ID      Pri      Out-
Ext-Color

P2-7009#show ip route vrf all database bgp
IP Route Table for VRF "default"

Total number of IPv4 routes 0
IP Route Table for VRF "management"

Total number of IPv4 routes 0

Gateway of last resort is not set
P2-7009#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF |      DESTINATION      | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref | Prefix |
| MPLS | OUT | OUT | NEXTHOP |
| ID |      FEC      | Ifname | IFNAME | LABEL |
| OBJECT | cnt | cnt |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
P2-7009#sh evpn mpls tunnel
% Evpn mpls is not Enabled !
P2-7009#

apps-fileview.texmex_20251030.00_p0
P2.txt
Displaying P2.txt.

```

PE3

```

PE3-7028#sh running-config
!
! Software version: HFCL-DUAR-OcNOS-SP-PLUS-7.0.0.136-Alpha
10/1
!
! Last configuration change at 13:45:22 UTC Fri Sep 12 2025 by
r
!
!
service password-encryption
!
logging console 5
logging logfile rsvp 7
logging level nsm 4
logging level rsvp 7
logging level bgp 7
logging level cmm 4
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile filter evpn-mpls-mh enable

```

```
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
qos enable
!
hostname PE3-7028
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
evpn mpls enable
!
evpn mpls irb
!
evpn mpls multihoming enable
!
ip vrf ip_vrfirb
  rd 28.28.28.28:200
  route-target both 200:200
  l3vni 20000
!
ip vrf ip_vrfirbMH801
  rd 28.28.28.28:801
  route-target both 801:801
  l3vni 40000
!
ip vrf management
!
mac vrf vrfirb
  rd 28.28.28.28:2000
  route-target both 2000:2000
!
mac vrf vrfirbMH801
  rd 28.28.28.28:4000
  route-target both 4000:4000
!
evpn mpls vtep-ip-global 28.28.28.28
!
evpn mpls id 102 xconnect target-mpls-id 101
!
evpn mpls id 201
  host-reachability-protocol evpn-bgp vrfirb
  evpn irb irb101
!
evpn mpls id 811
  host-reachability-protocol evpn-bgp vrfirbMH801
  evpn irb irb711
!
router rsvp
  auto-bypass
  attributes best-effort
  protection-capability node
  reoptimize
  exit
  inactivity-timer 5
  enable
  exit
```

```
!  
interface po20  
  switchport  
  evpn multi-homed system-mac 0000.1111.2222  
!  
interface po20.201 switchport  
  encapsulation dot1q 201  
  rewrite pop  
  access-if-evpn  
  map vpn-id 201  
!  
interface po20.811 switchport  
  encapsulation dot1q 811  
  rewrite pop  
  access-if-evpn  
  map vpn-id 811  
!  
interface ce1  
!  
interface ce2  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface ge25  
!  
interface irb101  
  ip vrf forwarding ip_vrfirb  
  ip address 90.90.1.1/24 anycast  
  ipv6 address 90::1/64 anycast  
!  
interface irb711  
  ip vrf forwarding ip_vrfirbMH801  
  evpn irb-if-forwarding anycast-gateway-mac  
  ip address 91.91.1.1/24 anycast  
  ipv6 address 91::1/64 anycast  
  ip ospf network point-to-point  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 28.28.28.28/32 secondary  
  ipv6 address ::1/128  
  ip router isis 1  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xe1  
!  
interface xe2  
!  
interface xe3  
!  
interface xe4  
!  
interface xe5  
!  
interface xe6  
!  
interface xe7  
!  
interface xe8  
!  
interface xe9  
!
```

```
interface xe10
!
interface xe11
load-interval 30
ip address 40.40.40.2/24
mtu 9216
label-switching
ip router isis 1
enable-rsvp
!
interface xe12
!
interface xe13
!
interface xe14
!
interface xe15
!
interface xe16
!
interface xe17
!
interface xe18
!
interface xe19
!
interface xe20
!
interface xe21
load-interval 30
ip address 30.30.40.2/24
mtu 9216
label-switching
ip router isis 1
enable-rsvp
!
interface xe22
!
interface xe23
!
interface xe24
speed 10g
channel-group 20 mode active
!
evpn irb-forwarding anycast-gateway-mac 0011.2200.0002
!
exit
!
router ospf 100
bfd all-interfaces
network 28.28.28.28/32 area 0.0.0.0
network 30.30.40.0/24 area 0.0.0.0
network 40.40.40.0/24 area 0.0.0.0
!
router bgp 65100
bgp router-id 28.28.28.28
no bgp inbound-route-filter
allocate-label all
neighbor 11.11.11.11 remote-as 65100
neighbor 11.11.11.11 update-source lo
neighbor 91.81.1.2 remote-as 65102
!
address-family ipv4 unicast
network 28.28.28.28/32
exit-address-family
!
address-family l2vpn evpn
neighbor 11.11.11.11 activate
```

```

neighbor 11.11.11.11 route-reflector-client
exit-address-family
!
address-family ipv4 vrf ip_vrfirb
redistribute connected
exit-address-family
!
address-family ipv4 vrf ip_vrfirbMH801
max-paths ibgp 4
redistribute connected
redistribute static
neighbor 91.91.1.2 remote-as 65102
neighbor 91.91.1.2 activate
exit-address-family
!
exit
!
rsvp-path PE3-PE1 mpls
30.30.40.1 strict
10.10.10.1 strict
!
rsvp-trunk PE3-PE1 ipv4
no primary cspf
primary path PE3-PE1
to 42.42.42.42
!
rsvp-trunk PE3-PE2 ipv4
to 30.30.30.30
!
rsvp-trunk PE3-PE4 ipv4
to 19.19.19.19
!
!
end
!

```

PE3-7028#sh ip ospf neighbor

Total number of full neighbors: 2

OSPF process 100 VRF(default):

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 11.11.11.11 | 1 | Full/Backup | 00:00:37 | 30.30.40.1 | xe21 | 0 |
| 9.9.9.9 | 1 | Full/Backup | 00:00:29 | 40.40.40.1 | xe11 | 0 |

PE3-7028#

PE3-7028#sh rsvp session

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass

State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary

* indicates the session is active with local repair at one or more nodes

(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|----------------------------------|-------------|--------|----------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 42.42.42.42 | 28.28.28.28 | 5001 | 2201 | PRI | PE3-PE1- | |
| Primary | | UP | 00:22:34 | 1 1 SE | - | 24330 |
| 30.30.30.30 | 28.28.28.28 | 5002 | 2202 | PRI | PE3-PE2- | |
| Primary | | UP | 00:27:28 | 1 1 SE | - | 24320 |
| 19.19.19.19 | 28.28.28.28 | 5003 | 2203 | PRI | PE3-PE4- | |
| Primary | | UP | 00:24:33 | 1 1 SE | - | 24324 |
| Total 3 displayed, Up 3, Down 0. | | | | | | |

Egress RSVP:

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|-------------|--------|---------|----------|----------|-------|
| Uptime | Rt | Style | Labelin | Labelout | | |
| 28.28.28.28 | 30.30.30.30 | 5002 | 2201 | PRI | PE2-PE3- | |

```

Primary          UP      00:27:22  1 1 SE      24960      -
28.28.28.28      42.42.42.42      5002      2203      PRI      PE1-PE3-
Primary          UP      00:24:01  1 1 SE      24962      -
28.28.28.28      19.19.19.19      5003      2203      PRI      PE4-PE3-
Primary          UP      00:24:57  1 1 SE      24961      -
Total 3 displayed, Up 3, Down 0.

```

```

PE3-7028#
PE3-7028#
PE3-7028#show bgp l2vpn evpn
BGP table version is 8, local router ID is 28.28.28.28
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

```

```
[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
```

```

1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

```

| Network | Next Hop | Metric | LocPrf | Weight | Path | Peer | Encap |
|---|-------------|--------|--------|--------|------|-------------|-------|
| RD[65100:501] | | | | | | | |
| *>i [3]:[501]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:502] | | | | | | | |
| *>i [3]:[502]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:503] | | | | | | | |
| *>i [3]:[503]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:504] | | | | | | | |
| *>i [3]:[504]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:505] | | | | | | | |
| *>i [3]:[505]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:506] | | | | | | | |
| *>i [3]:[506]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:507] | | | | | | | |
| *>i [3]:[507]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:508] | | | | | | | |
| *>i [3]:[508]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:509] | | | | | | | |
| *>i [3]:[509]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:510] | | | | | | | |
| *>i [3]:[510]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[19.19.19.19:250] | | | | | | | |
| *>i [5]:[0]:[0]:[24]:[70.70.1.0]:[0.0.0.0]:[16] | | | | | | | |


```

19.19.19.19      0      100      0      ?      - 11.11.11.11      MPLS

RD[19.19.19.19:801]
*>i      [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
19.19.19.19      0      100      0      ?      - 11.11.11.11      MPLS

RD[19.19.19.19:3000]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[251]:[19]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS
*>i      [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS
*>i      [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1]:[19]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS
*>i      [3]:[251]:[32,19.19.19]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS

RD[19.19.19.19:4000]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS
*>i      [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS
*>i      [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS
*>i      [3]:[811]:[32,19.19.19]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS

RD[19.19.19.19:64512]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS
*>i      [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS

RD[28.28.28.28:2000] VRF[vrfirb]
*>      [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
28.28.28.28      0      100      32768      i      - -----
MPLS
* i      [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0003]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0004]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0005]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0006]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0007]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0008]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0009]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:000a]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:000b]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:000c]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:000d]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:000e]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:000f]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS

```

[illegible]

[illegible]

```

* i [2]:[0]:[200]:[48,1200:0100:0054]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0055]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0056]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0057]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0058]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0059]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*> [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
    28.28.28.28 0 100 32768 i - -----
MPLS
*> [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1]:[19]
    28.28.28.28 0 100 32768 i - -----
MPLS
* i [3]:[200]:[32,42.42.42.42]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*> [3]:[201]:[32,28.28.28.28]
    28.28.28.28 0 100 32768 i - -----
MPLS

RD[28.28.28.28:4000] VRF[vrfirbMH801]
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
    19.19.19.19 0 100 0 i - 11.11.11.11 MPLS
*> 28.28.28.28 0 100 32768 i - -----
MPLS
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
    19.19.19.19 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1]:[19]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
    19.19.19.19 0 100 0 i - 11.11.11.11 MPLS
*> 28.28.28.28 0 100 32768 i - -----
MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
    19.19.19.19 0 100 0 i - 11.11.11.11 MPLS
*> 28.28.28.28 0 100 32768 i - -----
MPLS

```

```

* i   [3]:[801]:[32,42.42.42.42]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i   [3]:[811]:[32,19.19.19.19]
      19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS
*>   [3]:[811]:[32,28.28.28.28]
      28.28.28.28      0      100      32768      i      - -----
MPLS

RD[28.28.28.28:64512] VRF[evpn-gvrf-1]
*>   [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28      0      100      32768      i      - -----
MPLS
* i   [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
      19.19.19.19      0      100      0      i      - 11.11.11.11      MPLS
*>   [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28      0      100      32768      i      - -----
MPLS

RD[30.30.30.30:250]
*>i   [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
      30.30.30.30      0      100      0      ?      - 11.11.11.11      MPLS

RD[30.30.30.30:3000]
*>i   [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
      30.30.30.30      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1]:[19]
      30.30.30.30      0      100      0      i      - 11.11.11.11      MPLS
*>i   [3]:[250]:[32,30.30.30.30]
      30.30.30.30      0      100      0      i      - 11.11.11.11      MPLS

RD[42.42.42.42:801]
*>i   [5]:[0]:[0]:[24]:[81.81.1.0]:[0.0.0.0]:[17]
      42.42.42.42      0      100      0      ?      - 11.11.11.11      MPLS

RD[42.42.42.42:2000]
*>i   [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0003]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0004]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0005]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0006]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0007]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0008]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0009]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:000a]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:000b]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:000c]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:000d]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:000e]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:000f]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i   [2]:[0]:[200]:[48,1200:0100:0010]:[0]:[18]

```

2474

2475

```

42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0055]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0056]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0057]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0058]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0059]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [3]:[200]:[32,42.42.42.42]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS

RD[42.42.42.42:4000]
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1]:[19]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [3]:[801]:[32,42.42.42.42]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS

Total number of prefixes 254
PE3-7028#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code   FEC
Label  Out-Intf      Nexthop      UpTime      FTN-ID VRF-ID  Nhlfe-ID  Pri  Out-
B(x)>  81.81.1.0/24
      -          00:22:37 -          1          3          7          -   -
      42.42.42.42 -          -          5          Yes  17          -

PE3-7028#show ip route vrf all database bgp
IP Route Table for VRF "default"

Total number of IPv4 routes 0
IP Route Table for VRF "management"

```



```

Total number of IPv4 routes 0
IP Route Table for VRF "ip_vrfirb"

Total number of IPv4 routes 0
IP Route Table for VRF "ip_vrfirbMH801"
B      > 81.81.1.0/24 [200/0] via 42.42.42.42, installed 00:23:07, last update 00:23:07 ago

Total number of IPv4 routes 1
IP Route Table for VRF "vrfirb"

Total number of IPv4 routes 0
IP Route Table for VRF "vrfirbMH801"

Total number of IPv4 routes 0
IP Route Table for VRF "evpn-gvrf-1"

Total number of IPv4 routes 0

Gateway of last resort is not set
PE3-7028#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF |      DESTINATION      | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref | Prefix |
| MPLS |      OUT      | OUT      | NEXTHOP      |         |     |         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| ID |      FEC      | Ifname | IFNAME | LABEL |
| OBJECT | cnt | cnt |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
3    B> 81.81.1.0/24      7(s) - - -
      0x2000ccdc 00:22:37 1 1
      - /- /PRI /5 -
      17 42.42.42.42 0x2000ccdc

PE3-7028#sh evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      evpn-id      L
Leaf Remote-Leaf Ext-Color FAT
=====
28.28.28.28 42.42.42.42 Installed 00:23:09 00:23:09 40000
---
28.28.28.28 19.19.19.19 Installed 00:24:38 00:24:38 811
---
28.28.28.28 19.19.19.19 Installed 00:25:03 00:25:03 40000
---

Total number of entries are 3
PE3-7028#
apps-fileview.texmex_20251030.00_p0
PE3.txt
Displaying PE3.txt.

```

PE4

```
PE4-7019#sh running-config
!
! Software version: EC_AS5916-54X-OcNOS-SP-MPLS-7.0.0.136-
Alpha
!
! Last configuration change at 12:42:34 UTC Wed Oct 15 2025 by
r
!
!
```

```
service password-encryption
!
logging console 5
logging logfile rsvp 7
logging level nsm 4
logging level rsvp 7
logging level bgp 7
logging level cmm 4
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile filter evpn-mpls-mh enable
hardware-profile statistics ingress-acl enable
!
qos enable
!
hostname PE4-7019
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
evpn mpls enable
!
evpn mpls irb
!
evpn mpls multihoming enable
!
ip vrf ip_vrfirb250
  rd 19.19.19.19:250
  route-target both 250:250
  l3vni 30000
!
ip vrf ip_vrfirbMH801
  rd 19.19.19.19:801
  route-target both 801:801
  l3vni 40000
!
ip vrf management
!
mac vrf vrfirb250
  rd 19.19.19.19:3000
  route-target both 3000:3000
!
mac vrf vrfirbMH801
  rd 19.19.19.19:4000
  route-target both 4000:4000
!
evpn mpls vtep-ip-global 19.19.19.19
!
evpn mpls id 251
  host-reachability-protocol evpn-bgp vrfirb250
  evpn irb irb151
!
evpn mpls id 811
```

```
host-reachability-protocol evpn-bgp vrfirbMH801
evpn irb irb711
!
router rsvp
auto-bypass
attributes best-effort
protection-capability node
reoptimize
exit
inactivity-timer 5
enable
exit
!
interface po20
switchport
evpn multi-homed system-mac 0000.1111.2222
!
interface po20.251 switchport
encapsulation dot1q 251
rewrite pop
access-if-evpn
map vpn-id 251
!
interface po20.811 switchport
encapsulation dot1q 811
rewrite pop
access-if-evpn
map vpn-id 811
!
interface ce0
!
interface ce1
!
interface ce2
!
interface ce3
!
interface ce4
!
interface ce5
!
interface eth0
ip vrf forwarding management
ip address dhcp
!
interface irb151
ip vrf forwarding ip_vrfirb250
ip address 70.70.1.1/24 anycast
ipv6 address 70::1/64 anycast
!
interface irb711
ip vrf forwarding ip_vrfirbMH801
evpn irb-if-forwarding anycast-gateway-mac
ip address 91.91.1.1/24 anycast
ipv6 address 91::1/64 anycast
ip ospf network point-to-point
!
interface lo
ip address 127.0.0.1/8
ip address 19.19.19.19/32 secondary
ipv6 address ::1/128
ip router isis 1
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
```

```
interface xe0
!
interface xe1
!
interface xe2
!
interface xe3
  channel-group 20 mode active
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
!
interface xe10
!
interface xe11
!
interface xe12
!
interface xe13
!
interface xe14
!
interface xe15
!
interface xe16
!
interface xe17
!
interface xe18
!
interface xe19
!
interface xe20
!
interface xe21
!
interface xe22
!
interface xe23
!
interface xe24
!
interface xe25
!
interface xe26
!
interface xe27
!
interface xe28
!
interface xe29
!
interface xe30
!
interface xe31
!
interface xe32
!
interface xe33
```

```
!  
interface xe34  
!  
interface xe35  
!  
interface xe36  
!  
interface xe37  
!  
interface xe38  
!  
interface xe39  
!  
interface xe40  
!  
interface xe41  
!  
interface xe42  
!  
interface xe43  
!  
interface xe44  
  load-interval 30  
  ip address 30.30.30.2/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-rsvp  
!  
interface xe45  
!  
interface xe46  
  load-interval 30  
  ip address 40.40.50.2/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-rsvp  
!  
interface xe47  
!  
evpn irb-forwarding anycast-gateway-mac 0011.2200.0002  
!  
  exit  
!  
router ospf 100  
  bfd all-interfaces  
  network 19.19.19.19/32 area 0.0.0.0  
  network 30.30.30.0/24 area 0.0.0.0  
  network 40.40.50.0/24 area 0.0.0.0  
!  
router bgp 65100  
  bgp router-id 19.19.19.19  
  no bgp inbound-route-filter  
  neighbor 11.11.11.11 remote-as 65100  
  neighbor 11.11.11.11 update-source lo  
  neighbor 91.81.1.2 remote-as 65102  
  !  
  address-family l2vpn evpn  
  neighbor 11.11.11.11 activate  
  neighbor 11.11.11.11 route-reflector-client  
  exit-address-family  
  !  
  address-family ipv4 vrf ip_vrfirb250  
  redistribute connected  
  exit-address-family  
  !  
  address-family ipv4 vrf ip_vrfirbMH801
```

```

max-paths ibgp 4
redistribute connected
redistribute static
neighbor 91.91.1.2 remote-as 65102
neighbor 91.91.1.2 activate
exit-address-family
!
exit
!
rsvp-path PE4-PE1 mpls
30.30.30.1 strict
10.10.10.1 strict
!
rsvp-trunk PE4-PE1 ipv4
no primary cspf
primary path PE4-PE1
to 42.42.42.42
!
rsvp-trunk PE4-PE2 ipv4
to 30.30.30.30
!
rsvp-trunk PE4-PE3 ipv4
to 28.28.28.28
!

```

```
PE4-7019#sh ip ospf neighbor
```

```
Total number of full neighbors: 2
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 11.11.11.11 | 1 | Full/Backup | 00:00:37 | 30.30.30.1 | xe44 | 0 |
| 9.9.9.9 | 1 | Full/Backup | 00:00:36 | 40.40.50.1 | xe46 | 0 |

```
PE4-7019#
```

```
PE4-7019#sh rsvp session
```

```
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
```

```
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
```

```
* indicates the session is active with local repair at one or more nodes
```

```
(P) indicates the secondary-priority session is acting as primary
```

```
Ingress RSVP:
```

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|---------------|----------|--------|------|----------|-------|
| Uptime Rt | Style Labelin | Labelout | | | | |
| 42.42.42.42 | 19.19.19.19 | 5001 | 2201 | PRI | PE4-PE1- | |
| Primary | UP | 00:23:52 | 1 1 SE | - | 24328 | |
| 30.30.30.30 | 19.19.19.19 | 5002 | 2202 | PRI | PE4-PE2- | |
| Primary | UP | 00:25:41 | 1 1 SE | - | 24323 | |
| 28.28.28.28 | 19.19.19.19 | 5003 | 2203 | PRI | PE4-PE3- | |
| Primary | UP | 00:25:50 | 1 1 SE | - | 24322 | |

Total 3 displayed, Up 3, Down 0.

```
Egress RSVP:
```

| To | From | Tun-ID | LSP-ID | Type | LSPName | State |
|-------------|---------------|----------|--------|-------|----------|-------|
| Uptime Rt | Style Labelin | Labelout | | | | |
| 19.19.19.19 | 28.28.28.28 | 5003 | 2203 | PRI | PE3-PE4- | |
| Primary | UP | 00:25:26 | 1 1 SE | 24960 | - | |
| 19.19.19.19 | 42.42.42.42 | 5003 | 2204 | PRI | PE1-PE4- | |
| Primary | UP | 00:24:54 | 1 1 SE | 24961 | - | |

Total 2 displayed, Up 2, Down 0.

```
PE4-7019#
```

```
PE4-7019#
```

```
PE4-7019#show bgp l2vpn evpn
```

```
BGP table version is 6, local router ID is 19.19.19.19
```

```
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,
```

```
l - labeled, S Stale
```

Origin codes: i - IGP, e - EGP, ? - incomplete
 Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]

- 1 - Ethernet Auto-discovery Route
- 2 - MAC/IP Route
- 3 - Inclusive Multicast Route
- 4 - Ethernet Segment Route
- 5 - Prefix Route

| Network | Next Hop | Metric | LocPrf | Weight | Path | Peer | Encap |
|---|-------------|--------|--------|---------|------|-------------|-------|
| RD[65100:501] | | | | | | | |
| *>i [3]:[501]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:502] | | | | | | | |
| *>i [3]:[502]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:503] | | | | | | | |
| *>i [3]:[503]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:504] | | | | | | | |
| *>i [3]:[504]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:505] | | | | | | | |
| *>i [3]:[505]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:506] | | | | | | | |
| *>i [3]:[506]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:507] | | | | | | | |
| *>i [3]:[507]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:508] | | | | | | | |
| *>i [3]:[508]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:509] | | | | | | | |
| *>i [3]:[509]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:510] | | | | | | | |
| *>i [3]:[510]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[19.19.19.19:3000] VRF[vrfirb250] | | | | | | | |
| *> [1]:[00:00:00:11:11:22:22:00:00:00]:[251]:[19] | 19.19.19.19 | 0 | 100 | 32768 i | - | ----- | |
| MPLS | | | | | | | |
| * i [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| * i [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1]:[19] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| *> [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19] | 19.19.19.19 | 0 | 100 | 32768 i | - | ----- | |
| MPLS | | | | | | | |
| *> [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1]:[19] | 19.19.19.19 | 0 | 100 | 32768 i | - | ----- | |
| MPLS | | | | | | | |
| * i [3]:[250]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |

```

*> [3]:[251]:[32,19.19.19.19]
      19.19.19.19      0      100      32768 i      -      -----
      MPLS

RD[19.19.19.19:4000] VRF[vrfirbMH801]
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
*>      19.19.19.19      0      100      32768 i      - -----
      MPLS
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
      42.42.42.42      0      100      0 i      - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1]:[19]
      42.42.42.42      0      100      0 i      - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
      42.42.42.42      0      100      0 i      - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
      42.42.42.42      0      100      0 i      - 11.11.11.11 MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
*>      19.19.19.19      0      100      32768 i      - -----
      MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
*>      19.19.19.19      0      100      32768 i      - -----
      MPLS
* i [3]:[801]:[32,42.42.42.42]
      42.42.42.42      0      100      0 i      - 11.11.11.11 MPLS
*> [3]:[811]:[32,19.19.19.19]
      19.19.19.19      0      100      32768 i      - -----
      MPLS
* i [3]:[811]:[32,28.28.28.28]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS

RD[19.19.19.19:64512] VRF[evpn-gvrf-1]
*> [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      19.19.19.19      0      100      32768 i      - -----
      MPLS
*> [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
      19.19.19.19      0      100      32768 i      - -----
      MPLS
* i [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS

RD[28.28.28.28:200]
*>i [5]:[0]:[0]:[24]:[90.90.1.0]:[0.0.0.0]:[16]
      28.28.28.28      0      100      0 ?      - 11.11.11.11 MPLS

RD[28.28.28.28:801]
*>i [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
      28.28.28.28      0      100      0 ?      - 11.11.11.11 MPLS

RD[28.28.28.28:2000]
*>i [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
*>i [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
*>i [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1]:[19]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
*>i [3]:[201]:[32,28.28.28.28]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS

RD[28.28.28.28:4000]
*>i [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
*>i [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS

```



```

*>i    [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1][20]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS
*>i    [3]:[811]:[32,28.28.28]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS

RD[28.28.28.28:64512]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS
*>i    [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS

RD[30.30.30.30:250]
*>i    [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
      30.30.30.30      0      100      0      ?      - 11.11.11.11      MPLS

RD[30.30.30.30:3000]
*>i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
      30.30.30.30      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1][19]
      30.30.30.30      0      100      0      i      - 11.11.11.11      MPLS
*>i    [3]:[250]:[32,30.30.30.30]
      30.30.30.30      0      100      0      i      - 11.11.11.11      MPLS

RD[42.42.42.42:801]
*>i    [5]:[0]:[0]:[24]:[81.81.1.0]:[0.0.0.0]:[17]
      42.42.42.42      0      100      0      ?      - 11.11.11.11      MPLS

RD[42.42.42.42:2000]
*>i    [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0003]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0004]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0005]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0006]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0007]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0008]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0009]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:000a]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:000b]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:000c]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:000d]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:000e]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:000f]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0010]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0011]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0012]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0013]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS

```

2486

[illegible]

```

*>i [2]:[0]:[200]:[48,1200:0100:0058]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0059]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [3]:[200]:[32,42.42.42.42]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS

```

RD[42.42.42.42:4000]

```

*>i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1]:[19]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [3]:[801]:[32,42.42.42.42]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS

```

Total number of prefixes 155

PE4-7019#show mpls vrf-forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN

(m) - Service mapped over multipath transport

(e) - Service mapped over ECMP

(D) - Down

Ext-Color - Extended-community color advertised by BGP

B(x) - BGP EVPN MPLS Services

| Code | FEC | Nexthop | UpTime | FTN-ID | VRF-ID | Nhlfe-ID | Pri | Out- |
|-------|-----------------|----------|--------|-----------|--------|----------|-----|------|
| Label | Out-Intf | | | Ext-Color | | | | |
| B(x)> | 60.60.1.0/24 | | | 1 | 2 | 4 | - | - |
| | - | 00:25:57 | - | | | 3 | Yes | 16 |
| e44 | (D) 30.30.30.30 | - | - | | | | | x |
| B(x)> | 81.81.1.0/24 | | | 1 | 3 | 9 | - | - |
| | - | 00:23:54 | - | | | 7 | Yes | 17 |
| | 42.42.42.42 | - | - | | | | | - |

PE4-7019#show ip route vrf all database bgp

IP Route Table for VRF "default"

Total number of IPv4 routes 0

IP Route Table for VRF "management"

Total number of IPv4 routes 0

IP Route Table for VRF "ip_vrfirb250"

```
B > 60.60.1.0/24 [200/0] via 30.30.30.30, installed 00:25:57, last update 00:25:57 ago
```

```
Total number of IPv4 routes 1
```

```
IP Route Table for VRF "ip_vrfirbMH801"
```

```
B > 81.81.1.0/24 [200/0] via 42.42.42.42, installed 00:24:00, last update 00:24:00 ago
```

```
Total number of IPv4 routes 1
```

```
IP Route Table for VRF "vrfirb250"
```

```
Total number of IPv4 routes 0
```

```
IP Route Table for VRF "vrfirbMH801"
```

```
Total number of IPv4 routes 0
```

```
IP Route Table for VRF "evpn-gvrf-1"
```

```
Total number of IPv4 routes 0
```

```
Gateway of last resort is not set
```

```
PE4-7019#show hsl mpls l3vpn-ftn
```

```
TABLE: L3VPN Ftn table
```

```
Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup
```

```
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy
```

```
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF | DESTINATION | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref | Prefix |
| MPLS | OUT | OUT | NEXTHOP | | | |
| ID | FEC | cnt | cnt | Ifname | IFNAME | LABEL |
| OBJECT | | | | | | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
3 B> 81.81.1.0/24 9(s) - - - -
0x2000101a 00:23:54 1 1
- /- /PRI /7 -
17 42.42.42.42 0x2000101a
```

```
PE4-7019#sh evpn mpls tunnel
```

```
EVPN-MPLS Network tunnel Entries
```

| Source | Destination | Status | Up/Down | Update | evpn-id | Local- |
|-------------|-------------|-----------|----------|----------|---------|--------|
| Leaf | Remote-Leaf | Ext-Color | FAT | | | |
| 19.19.19.19 | 28.28.28.28 | Installed | 00:25:54 | 00:25:54 | 811 | --- |
| 19.19.19.19 | 28.28.28.28 | Installed | 00:25:59 | 00:25:59 | 40000 | --- |
| 19.19.19.19 | 42.42.42.42 | Installed | 00:24:02 | 00:24:02 | 40000 | --- |
| 19.19.19.19 | 30.30.30.30 | Installed | 00:25:59 | 00:25:59 | 30000 | --- |

```
Total number of entries are 4
```

```
PE4-7019#
```

```
apps-fileview.texmex_20251030.00_p0
```

```
PE4.txt
```

```
Displaying PE4.txt.
```

Validation

To verify the ECMP configuration, check the output of the following `show` commands:

```
PE4-7019#show bgp l2vpn evpn
```

```
BGP table version is 6, local router ID is 19.19.19.19
```

```
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,
```

```
l - labeled, S Stale
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
Description : Ext-Color - Extended community color
```

```
[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
```

- 1 - Ethernet Auto-discovery Route
- 2 - MAC/IP Route
- 3 - Inclusive Multicast Route
- 4 - Ethernet Segment Route
- 5 - Prefix Route

| Network | Next Hop | Metric | LocPrf | Weight | Path | Peer | Encap |
|---|-------------|--------|--------|---------|------|-------------|-------|
| RD[65100:501] | | | | | | | |
| *>i [3]:[501]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:502] | | | | | | | |
| *>i [3]:[502]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:503] | | | | | | | |
| *>i [3]:[503]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:504] | | | | | | | |
| *>i [3]:[504]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:505] | | | | | | | |
| *>i [3]:[505]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:506] | | | | | | | |
| *>i [3]:[506]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:507] | | | | | | | |
| *>i [3]:[507]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:508] | | | | | | | |
| *>i [3]:[508]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:509] | | | | | | | |
| *>i [3]:[509]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[65100:510] | | | | | | | |
| *>i [3]:[510]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| RD[19.19.19.19:3000] VRF[vrfirb250] | | | | | | | |
| *> [1]:[00:00:00:11:11:22:22:00:00]:[251]:[19] | 19.19.19.19 | 0 | 100 | 32768 i | - | ----- | |
| MPLS | | | | | | | |
| * i [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| * i [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1]:[19] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| *> [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19] | 19.19.19.19 | 0 | 100 | 32768 i | - | ----- | |
| MPLS | | | | | | | |
| *> [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1]:[19] | 19.19.19.19 | 0 | 100 | 32768 i | - | ----- | |
| MPLS | | | | | | | |
| * i [3]:[250]:[32,30.30.30.30] | 30.30.30.30 | 0 | 100 | 0 i | - | 11.11.11.11 | MPLS |
| *> [3]:[251]:[32,19.19.19.19] | 19.19.19.19 | 0 | 100 | 32768 i | - | ----- | |
| MPLS | | | | | | | |

```

RD[19.19.19.19:4000] VRF[vrfirbMH801]
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
*> 19.19.19.19 0 100 32768 i - -----
      MPLS
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1][19]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1][19]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
*> 19.19.19.19 0 100 32768 i - -----
      MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1][20]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
*> 19.19.19.19 0 100 32768 i - -----
      MPLS
* i [3]:[801]:[32,42.42.42.42]
      42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*> [3]:[811]:[32,19.19.19.19]
      19.19.19.19 0 100 32768 i - -----
      MPLS
* i [3]:[811]:[32,28.28.28.28]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS

RD[19.19.19.19:64512] VRF[evpn-gvrf-1]
*> [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      19.19.19.19 0 100 32768 i - -----
      MPLS
*> [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
      19.19.19.19 0 100 32768 i - -----
      MPLS
* i [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS

RD[28.28.28.28:200]
*>i [5]:[0]:[0]:[24]:[90.90.1.0]:[0.0.0.0]:[16]
      28.28.28.28 0 100 0 ? - 11.11.11.11 MPLS

RD[28.28.28.28:801]
*>i [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
      28.28.28.28 0 100 0 ? - 11.11.11.11 MPLS

RD[28.28.28.28:2000]
*>i [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1][19]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
*>i [3]:[201]:[32,28.28.28.28]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS

RD[28.28.28.28:4000]
*>i [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1][20]
      28.28.28.28 0 100 0 i - 11.11.11.11 MPLS
*>i [3]:[811]:[32,28.28.28.28]

```

```

28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS

RD[28.28.28.28:64512]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS
*>i      [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS

RD[30.30.30.30:250]
*>i      [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
30.30.30.30      0      100      0      ?      -  11.11.11.11      MPLS

RD[30.30.30.30:3000]
*>i      [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1][19]
30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS
*>i      [3]:[250]:[32,30.30.30.30]
30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS

RD[42.42.42.42:801]
*>i      [5]:[0]:[0]:[24]:[81.81.1.0]:[0.0.0.0]:[17]
42.42.42.42      0      100      0      ?      -  11.11.11.11      MPLS

RD[42.42.42.42:2000]
*>i      [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0003]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0004]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0005]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0006]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0007]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0008]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0009]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:000a]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:000b]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:000c]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:000d]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:000e]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:000f]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0010]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0011]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0012]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0013]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0014]:[0]:[18]
42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i      [2]:[0]:[200]:[48,1200:0100:0015]:[0]:[18]

```


[illegible]

© 2026 IP Infusion Inc. Proprietary 2494

```

42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [3]:[200]:[32,42.42.42]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS

RD[42.42.42.42:4000]
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1][19]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1][19]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i [3]:[801]:[32,42.42.42]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS

Total number of prefixes 155
PE4-7019#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code   FEC
Label  Out-Intf      Nexthop      UpTime      FTN-ID VRF-ID  Nhlfe-ID  Pri  Out-
B(x)>  60.60.1.0/24      -      00:25:57  -      1      2      4      -  -  -
                                     3      Yes  16  x
e44    (D) 30.30.30.30      -      -
B(x)>  81.81.1.0/24      -      00:23:54  -      1      3      9      -  -  -
                                     7      Yes  17  -
                                     42.42.42.42      -      -

PE4-7019#show ip route vrf all database bgp
IP Route Table for VRF "default"

Total number of IPv4 routes 0
IP Route Table for VRF "management"

Total number of IPv4 routes 0
IP Route Table for VRF "ip_vrfrb250"
B      > 60.60.1.0/24 [200/0] via 30.30.30.30, installed 00:25:57, last update 00:25:57 ago

Total number of IPv4 routes 1

```

```

IP Route Table for VRF "ip_vrfirbMH801"
B    > 81.81.1.0/24 [200/0] via 42.42.42.42, installed 00:24:00, last update 00:24:00 ago

Total number of IPv4 routes 1
IP Route Table for VRF "vrfirb250"

Total number of IPv4 routes 0
IP Route Table for VRF "vrfirbMH801"

Total number of IPv4 routes 0
IP Route Table for VRF "evpn-gvrf-1"

Total number of IPv4 routes 0

Gateway of last resort is not set
PE4-7019#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF | DESTINATION | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref | Prefix |
| MPLS | OUT | OUT | NEXTHOP | | | |
| ID | FEC | Ifname | IFNAME | LABEL |
| OBJECT | cnt | cnt |
+-----+-----+-----+-----+-----+-----+-----+-----+
3    B> 81.81.1.0/24          9(s) -          -          -          -
      0x2000101a 00:23:54  1    1
      -    /-    /PRI    /7    -
      17      42.42.42.42    0x2000101a

PE4-7019#sh evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination    Status      Up/Down      Update      evpn-id      Local-
Leaf Remote-Leaf Ext-Color  FAT
=====
19.19.19.19  28.28.28.28    Installed   00:25:54     00:25:54     811          ---
---          ---
19.19.19.19  28.28.28.28    Installed   00:25:59     00:25:59     40000        ---
---          ---
19.19.19.19  42.42.42.42    Installed   00:24:02     00:24:02     40000        ---
---          ---
19.19.19.19  30.30.30.30    Installed   00:25:59     00:25:59     30000        ---
---          ---

Total number of entries are 4
PE4-7019#

```

ECMP for L3EVPN Commands

The ECMP feature in L3EVPN has the following configuration commands:



Note: The `multipath` CLI is available under the EVPN address family. When this command is configured, BGP computes multiple best paths per Route Distinguisher (RD) and advertises these routes to the associated VRFs. The VRFs then install the routes based on the `max-paths` configuration defined at the VRF level. If the `multipath` CLI is not configured and only `max-paths` is configured under the VRF, BGP propagates only the single best route from the RD to the VRF. The VRF then installs multiple paths based on the `max-paths` configuration.

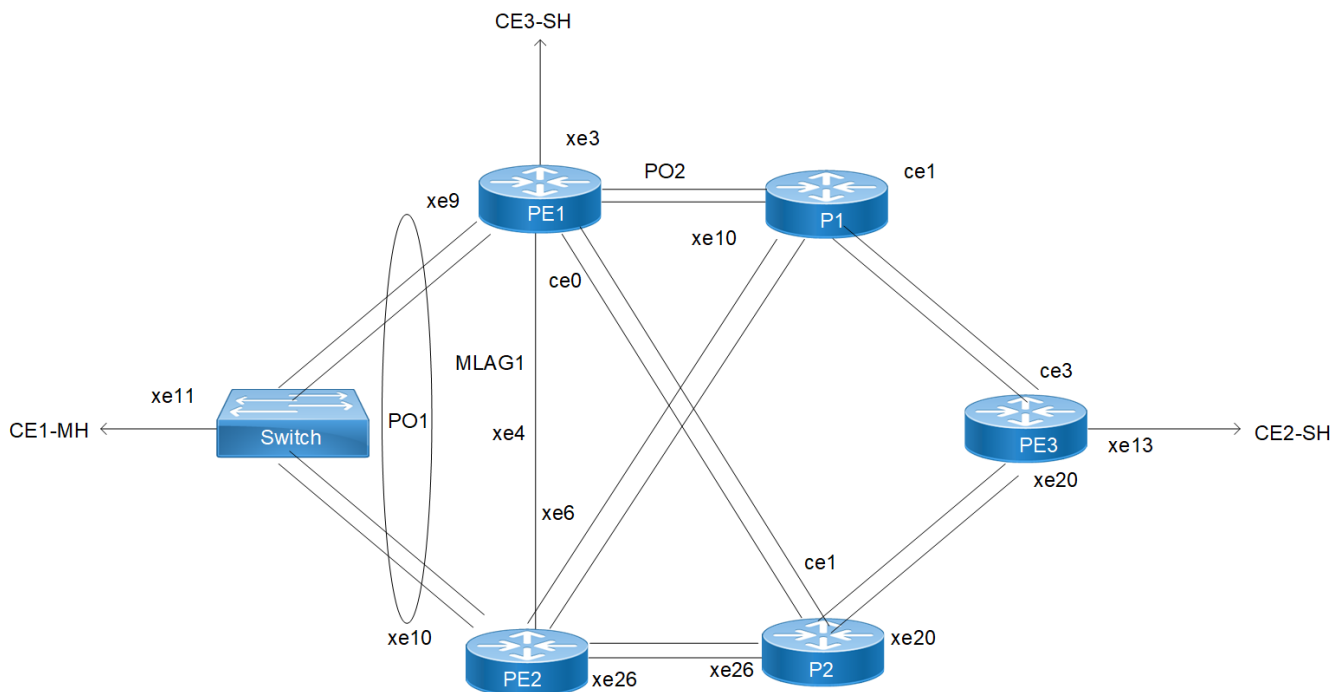
L3VPN over EVPN MPLS Configuration

L3vpn EVPN MPLS ensures that we are able to do the Ip-vrf to Ip-vrf routing with MPLS as overlay. We can ensure that with this model we can do the Anycast gateway using EVPN concepts. L3VPN EVPN MPLS is a way of integrating the interface-less model. In this model we will use non-irb interface i.e. L3 interface for the Ip-Vrf routing. Using Anycast gateway idea is to have the Multi-Homed PE nodes to work in active-standby mode. From remote perspective at one time only one Peer PE will be active. If that PE goes down, then a new tunnel to the peer PE will be established for the traffic. From the access-side Traffic can reach to either of the peer PE and sent across the remote PE nodes. This way we can achieve redundancy using the anycast mac-address.

Topology

Below Topology depicts the topology for the L3VPN over EVPN MPLS configuration examples for both single homing and multi-homing with SR/LDP. MLAG configured between PE1 and PE2 to achieve multi-homing Active-Standby connected to CE1.

Figure 97. Topology diagram for L3VPN over EVPN MPLS



Configuration

PE1

Loopback Interface:

```
#configure terminal
```

Enter configuration mode.

| | |
|--|--|
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)# ip address 1.1.1.1/32 secondary | Configure IP address on loopback interface. |
| (config-if)# prefix-sid index 1 | Configuring prefix sid for segment id |
| (config-if)#commit | Commit the transaction. |
| (config-if)#exit | Exit interface mode |

Global EVPN MPLS Command:

| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#load-balance enable | Enable Load balance |
| (config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware profile for evpn mpls multihoming |
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#evpn mpls irb | Enable EVPN MPLS IRB |
| (config)#evpn mpls multihoming enable | Enable EVPN MPLS multihoming |
| (config)#evpn mpls vtep-ip-global 1.1.1.1 | Configuring VTEP global IP to loopback IP |
| (config)#evpn irb-forwarding anycast-gateway-mac 0011.2233.4466 | Configuring anycast gateway mac for MH. Note: Anycast gateway is mandatory for MH nodes. |
| (config)#commit | Commit candidate configuration to be running configuration Note: Reload is required after Enabling/Disabling EVPN MPLS Feature. |
| (config)#segment-routing | Configure segment routing |
| (config-sr)# mpls sr-prefer | Set mpls prefer segment routing over other protocols |
| (config-sr)#exit | Exit from router sr mode |
| (config)#commit | Commit the transaction. |

Configure LDP:

| | |
|---|---|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)# transport-address ipv4 1.1.1.1 | Configuring transport address |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit the transaction. |

MLAG Configuration:

| | |
|-------------------------------|--------------------------|
| (config)#interface mlag1 | Create mlag interface |
| (config-if)# switchport | Configuring as L2 port |
| (config-if)# load-interval 30 | Configure load interval |
| (config-if)# exit | Exit from interface mode |

| | |
|---|--|
| (config)#interface po1 | Configure dynamic lag |
| (config-if)# switchport | Configuring as L2 port |
| (config-if)# mtu 1500 | Set mtu value |
| (config-if)# mlag 1 | Attach mlag to po1. Note: While creating po1 subifp, this needs to be un-configured and configured back |
| (config-if)#interface xe10 | Enter interface mode for xe10 |
| (config-if)# channel-group 1 mode active | Attach lag interface po1 |
| (config-if)#exit | Exit interface mode |
| (config)#mcec domain configuration | Create mcec domain |
| (config-mcec-domain)# domain-address 1111.2222.3333 | Configure mcec domain address |
| (config-mcec-domain)# domain-system-number 1 | Configure system number. Active node should have lower value. |
| (config-mcec-domain)# intra-domain-link xe4 | Configure ideal interface between mlag devices |
| (config-mcec-domain)# domain-hello-timeout long | Configure domain hello timeout |

Interface Configuration Network Side:

| | |
|---|--|
| (config-if)#interface ce0 | Configure network interface ce0 |
| (config-if)# ip address 20.1.1.1/24 | Configure IP address on the interface. |
| (config-if)# mtu 1522 | Configure mtu |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)# mpls ldp-igp sync ospf | Configure mpls igp sync |
| (config-if)# ip ospf network point-to-point | Configure ospf as p2p |
| (config-if)# enable-ldp ipv4 | Enable LDP on the physical interface |
| (config)#interface po2 | Create channel group po2 |
| (config-if)# ip address 10.1.1.1/24 | Configure IP address on the interface. |
| (config-if)# mtu 2000 | Configure mtu |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)# mpls ldp-igp sync ospf | Configure mpls igp sync |
| (config-if)# ip ospf network point-to-point | Configure ospf as p2p |
| (config-if)# enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#interface xe18 | Enter interface mode for xe18 |
| (config-if)# channel-group 2 mode active | Attach lag interface po2 |
| (config-if)#interface xe19 | Enter interface mode for xe19 |
| (config-if)# channel-group 2 mode active | Attach lag interface po2 |
| (config)#commit | Commit the transaction. |

OSPF Configuration:

| | |
|--|---|
| (config)#router ospf 1 | Enter the Router OSPF mode. |
| (config-router)# network 1.1.1.1/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)# network 10.1.1.0/24 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# network 20.1.1.0/24 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# ospf segment-routing global block 16000 23000 | Configure SR global block for ospf |
| (config-router)# segment-routing mpls | Enable ospf SR |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit the transaction. |

BGP Configuration:

| | |
|--|---|
| (config)#router bgp 100 | Enter the Router BGP mode, ASN: 100 |
| (config-router)# neighbor 5.5.5.5 remote-as 100 | Configuring PE3 as iBGP neighbor using it's loopback IP |
| (config-router)# neighbor 5.5.5.5 fall-over bfd multihop | Configure BFD |
| (config-router)# neighbor 5.5.5.5 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 unicast | Enter in to address family ipv4 unicast mode |
| (config-router-af)# neighbor 5.5.5.5 activate | Activate PE3 neighbor |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family l2vpn evpn | Enter in to address family l2vpn vpn |
| (config-router-af)# neighbor 5.5.5.5 activate | Activate PE3 neighbor |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv6 unicast | Enter in to address family ipv6 unicast mode |
| (config-router-af)# neighbor 5.5.5.5 activate | Activate PE3 neighbor |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv4 vrf vrf102 | Enter in to address family ipv4 vrf vrf102 |
| (config-router-af)# redistribute connected | Redistribute connected networks in to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv4 vrf vrf101 | Enter in to address family ipv4 vrf vrf102 |
| (config-router-af)# redistribute connected | Redistribute connected networks in to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv6 vrf vrf101 | Enter in to address family ipv6 vrf vrf101 |
| (config-router-af)# redistribute connected | Redistribute connected networks in to bgp |
| (config-router-af)# exit-address-family | Exit address family |

| | |
|---|--|
| (config-router)# address-family ipv6 vrf vrf102 | Enter in to address family ipv6 vrf vrf102 |
| (config-router-af)# redistribute connected | Redistribute connected networks in to bgp |
| (config-router-af)# exit-address-family | Exit address family |

IP VRF Configuration:

| | |
|---|---|
| (config-vrf)#ip vrf vrf102 | Enter VRF mode |
| (config-vrf)# rd 10001:102 | Configuring Route-Distinguisher value |
| (config-vrf)# route-target both 1.1.1.1:102 | Configure RT value |
| (config-vrf)# l3vni 102 | Configure L3VNID to populate route through evpn |
| (config-vrf)#ip vrf vrf101 | Enter VRF mode |
| (config-vrf)# rd 10001:1 | Configuring Route-Distinguisher value |
| (config-vrf)# route-target both 1.1.1.1:1 | Configure RT value |
| (config-vrf)# l3vni 101 | Configure L3VNID to populate route through evpn |

L3 VRF Port Configuration:

| | |
|--|--|
| (config-if)#interface xe3 | Enter the Interface mode for xe3. This is for untagged traffic |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# ip address 100.1.1.1/24 | Assign ipv4 address |
| (config-if)# ipv6 address 1000::1/64 | Assign ipv6 address |
| (config-if)#interface xe3.1 | Create xe3.1 subifp |
| (config-if)# encapsulation dot1q 10 | Configure encapsulation as single tagged |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# ip address 110.1.1.1/24 | Assign ipv4 address |
| (config-if)# ipv6 address 1100::1/64 | Assign ipv6 address |
| (config-if)#interface xe3.2 | Create xe3.2 subifp |
| (config-if)# encapsulation dot1q 11 inner-dot1q 11 | Configure encapsulation as double tagged |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# ip address 120.1.1.1/24 | Assign ipv4 address |
| (config-if)# ipv6 address 1200::1/64 | Assign ipv6 address |
| (config-if)#interface xe3.11 | Create xe3.11 subifp |
| (config-if)# encapsulation dot1q 11 | Configure encapsulation as single tagged |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# ip address 111.1.1.1/24 | Assign ipv4 address |
| (config-if)# ipv6 address 1110::1/64 | Assign ipv6 address |
| (config-if)#interface po1.11 | Create po1 subinterface. We need to unconfigure mlag under po1 before creating subinterface. |

| | |
|---|--|
| (config-if)# encapsulation dot1q 11 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# evpn irb-if-forwarding anycast-gateway-mac | Configure anycast gateway mac |
| (config-if)# ip address 150.11.1.1/24 | Assign ipv4 address |
| (config-if)#interface po1.12 | Create po1 subinterface. |
| (config-if)# encapsulation dot1q 12 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# evpn irb-if-forwarding anycast-gateway-mac | Configure anycast gateway mac |
| (config-if)# ip address 150.12.1.1/24 | Assign ipv4 address |
| (config-if)#interface po1.13 | Create po1 subinterface. |
| (config-if)# encapsulation dot1q 13 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# evpn irb-if-forwarding anycast-gateway-mac | Configure anycast gateway mac |
| (config-if)# ip address 150.13.1.1/24 | Assign ipv4 address |
| (config-if)#interface po1.14 | Create po1 subinterface. |
| (config-if)# encapsulation dot1q 14 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# evpn irb-if-forwarding anycast-gateway-mac | Configure anycast gateway mac |
| (config-if)# ip address 150.14.1.1/24 | Assign ipv4 address |
| (config-if)#interface po1.15 | Create po1 subinterface. |
| (config-if)# encapsulation dot1q 15 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# evpn irb-if-forwarding anycast-gateway-mac | Configure anycast gateway mac |
| (config-if)# ip address 150.15.1.1/24 | Assign ipv4 address |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |
| PE2 Configuration: | |
| Loopback Interface: | |
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)# ip address 2.2.2.2/32 secondary | Configure IP address on loopback interface. |
| (config-if)# prefix-sid index 2 | Configuring prefix sid for segment id |
| (config-if)#commit | Commit the transaction. |
| (config-if)#exit | Exit interface mode |

Global EVPN MPLS Command:

| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#load-balance enable | Enable Load balance |
| (config)#hardware-profile filter evpn-mpls-mh enable | Enable hardware profile for evpn mpls multihoming |
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#evpn mpls irb | Enable EVPN MPLS IRB |
| (config)#evpn mpls multihoming enable | Enable EVPN MPLS multihoming |
| (config)#evpn mpls vtep-ip-global 2.2.2.2 | Configuring VTEP global IP to loopback IP |
| (config)#evpn irb-forwarding anycast-gateway-mac 0011.2233.4466 | Configure Anycast gateway mac for MH. This command is mandatory for MH nodes. |
| (config)#commit | Commit candidate configuration to be running configuration Note: Reload is required after Enabling/Disabling EVPN MPLS Feature. |

Configure SR:

| | |
|-----------------------------|--|
| (config)#segment-routing | Configure segment routing |
| (config-sr)# mpls sr-prefer | Set mpls prefer segment routing over other protocols |
| (config-sr)#exit | Exit from router sr mode |
| (config)#commit | Commit the transaction. |

Configure LDP:

| | |
|---|---|
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)# transport-address ipv4 2.2.2.2 | Configuring transport address |
| (config-router)#exit | Exit from router target peer and LDP mode |
| (config)#commit | Commit the transaction. |

MLAG Configuration:

| | |
|-------------------------------|--------------------------|
| (config)#interface mlag1 | Create mlag interface |
| (config-if)# switchport | Configuring as L2 port |
| (config-if)# load-interval 30 | Configure load interval |
| (config-if)# exit | Exit from interface mode |
| (config)#interface pol | Configure dynamic lag |
| (config-if)# switchport | Configuring as L2 port |
| (config-if)# mtu 1500 | Set mtu value |
| (config-if)# mlag 1 | Attach mlag |

| | |
|---|--|
| (config-if)#interface xe9 | Enter interface mode |
| (config-if)# speed 10g | Set speed as 10g |
| (config-if)# channel-group 1 mode active | Attach the channel group po1 |
| (config-if)# exit | Exit from interface mode |
| (config)#mcec domain configuration | Create mcec domain |
| (config-mcec-domain)# domain-address 1111.2222.3333 | Configure mcec domain address |
| (config-mcec-domain)# domain-system-number 2 | Configure system number |
| (config-mcec-domain)# intra-domain-link xe4 | Configure ideal interface between mlag devices |
| (config-mcec-domain)# domain-hello-timeout long | Configure domain hello timeout |
| (config)#commit | Commit the transaction. |

Interface Configuration Network Side:

| | |
|---|--|
| (config-if)#interface xe6 | Enter the Interface mode for xe6. |
| (config-if)# ip address 30.1.1.1/24 | Configure ipv4 address |
| (config-if)# mtu 1522 | Configure mtu |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf as p2p |
| (config-if)# enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#interface xe26 | Enter the Interface mode for xe26. |
| (config-if)# ip address 40.1.1.1/24 | Configure ipv4 address |
| (config-if)# mtu 2000 | Configure mtu |
| (config-if)# label-switching | Enable label switching on the interface. |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp sync with ospf |
| (config-if)# ip ospf network point-to-point | Configure ospf as p2p |
| (config-if)# enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-router)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |

OSPF Configuration:

| | |
|--|-------------------------------------|
| (config)#router ospf 1 | Enter the Router OSPF mode. |
| (config-router)# network 1.1.1.1/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)# network 10.1.1.0/24 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# network 20.1.1.0/24 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# ospf segment-routing global block 16000 23000 | Configure SR global block for ospf |
| (config-router)# segment-routing mpls | Enable ospf SR |

| | |
|----------------------|---|
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit the transaction. |

BGP Configuration:

| | |
|--|---|
| (config)#router bgp 100 | Enter the Router BGP mode, ASN: 100 |
| (config-router)# neighbor 5.5.5.5 remote-as 100 | Configuring PE3 as iBGP neighbor using it's loopback IP |
| (config-router)# neighbor 5.5.5.5 fall-over bfd multihop | Configure BFD |
| (config-router)# neighbor 5.5.5.5 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 unicast | Enter in to address family ipv4 unicast mode |
| (config-router-af)# neighbor 5.5.5.5 activate | Activate PE3 neighbor |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family l2vpn evpn | Enter in to address family l2vpn vpn |
| (config-router-af)# neighbor 5.5.5.5 activate | Activate PE3 neighbor |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv6 unicast | Enter in to address family ipv6 unicast mode |
| (config-router-af)# neighbor 5.5.5.5 activate | Activate PE3 neighbor |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv4 vrf vrf102 | Enter in to address family ipv4 vrf vrf102 |
| (config-router-af)# redistribute connected | Redistribute connected networks in to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv4 vrf vrf101 | Enter in to address family ipv4 vrf vrf101 |
| (config-router-af)# redistribute connected | Redistribute connected networks in to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv6 vrf vrf101 | Enter in to address family ipv6 vrf vrf101 |
| (config-router-af)# redistribute connected | Redistribute connected networks in to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv6 vrf vrf102 | Enter in to address family ipv6 vrf vrf102 |
| (config-router-af)# redistribute connected | Redistribute connected networks in to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config)#commit | Commit the transaction. |

IP VRF Configuration:

| | |
|----------------------------|---------------------------------------|
| (config-vrf)#ip vrf vrf102 | Enter VRF mode |
| (config-vrf)# rd 10001:102 | Configuring Route-Distinguisher value |

| | |
|---|---|
| (config-vrf)# route-target both 1.1.1.1:102 | Configure RT |
| (config-vrf)# l3vni 102 | Configure L3VNI to send ipv4 route through evpn |
| (config-vrf)#ip vrf vrf101 | Create ip vrf vrf101 |
| (config-vrf)# rd 10001:1 | Configuring Route-Distinguisher value |
| (config-vrf)# route-target both 1.1.1.1:1 | Configure RT |
| (config-vrf)# l3vni 101 | Configure L3VNI to send ipv4 route through evpn |
| (config)#commit | Commit the transaction. |

L3 VRF Port Configuration:

| | |
|---|--|
| (config-if)#interface po1.11 | Create po1 subinterface. We need to unconfigure mlag under po1 before creating subinterface. |
| (config-if)# encapsulation dot1q 11 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# evpn irb-if-forwarding anycast-gateway-mac | Configure anycast gateway mac for multi-homing |
| (config-if)# ip address 150.11.1.1/24 | Assign ipv4 address |
| (config-if)#interface po1.12 | Create po1 subinterface. |
| (config-if)# encapsulation dot1q 12 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# evpn irb-if-forwarding anycast-gateway-mac | Configure anycast gateway mac for multi-homing |
| (config-if)# ip address 150.12.1.1/24 | Assign ipv4 address |
| (config-if)#interface po1.13 | Create po1 subinterface. |
| (config-if)# encapsulation dot1q 13 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# evpn irb-if-forwarding anycast-gateway-mac | Configure anycast gateway mac for multi-homing |
| (config-if)# ip address 150.13.1.1/24 | Assign ipv4 address |
| (config-if)#interface po1.14 | Create po1 subinterface. |
| (config-if)# encapsulation dot1q 14 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# ip address 150.14.1.1/24 | Assign ipv4 address |
| (config-if)# evpn irb-if-forwarding anycast-gateway-mac | Configure anycast gateway mac for multi-homing |
| (config-if)#interface po1.15 | Create po1 subinterface. |
| (config-if)# encapsulation dot1q 15 | Configure encapsulation type and value |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# ip address 150.15.1.1/24 | Assign ipv4 address |
| (config-if)# evpn irb-if-forwarding anycast- | Configure anycast gateway mac for multi-homing |

| | |
|------------------|-------------------------|
| gateway-mac | |
| (config-if)#exit | Exit interface mode |
| (config)#commit | Commit the transaction. |

P1:

| | |
|---|----------------------------------|
| (config)#segment-routing | Configure SR. |
| (config-sr)# mpls sr-prefer | Prefer SR over other protocol. |
| (config-sr)#exit | Exit sr mode. |
| (config)#router ldp | Configure router ldp |
| (config-router)# transport-address ipv4 3.3.3.3 | Configure transport address. |
| (config-router)#exit | Exit |
| (config)#interface po2 | Configure dynamic LAG po2 |
| (config-if)# ip address 10.1.1.2/24 | Configure ipv4 address |
| (config-if)# mtu 2000 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |
| (config-if)#interface ce1 | Configure network interface ce1 |
| (config-if)# ip address 60.1.1.1/24 | Configure ipv4 address |
| (config-if)# mtu 1522 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |
| (config-if)#interface lo | Configure loopback interface |
| (config-if)# ip address 3.3.3.3/32 secondary | Configure secondary ip |
| (config-if)# prefix-sid absolute 16003 | Configure SR segment id |
| (config-if)#interface xe10 | Configure network interface xe10 |
| (config-if)# speed 10g | Set speed 10g |
| (config-if)# ip address 30.1.1.2/24 | Configure ipv4 address |
| (config-if)# mtu 1522 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |

| | |
|---|----------------------------------|
| (config-if)#interface xe14 | Configure network interface xe14 |
| (config-if)# ip address 50.1.1.1/24 | Configure ipv4 address |
| (config-if)# mtu 1522 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |
| (config-if)#interface xe16 | Enter interface mode |
| (config-if)# channel-group 2 mode active | Map dynamic lag po2 |
| (config-if)#interface xe17 | Enter interface mode |
| (config-if)# channel-group 2 mode active | Map dynamic lag po2 |
| (config-if)# exit | Exit |
| (config)#router ospf 1 | Configure router ospf 1 |
| (config-router)# network 3.3.3.3/32 area 0.0.0.0 | Add loopback ip to ospf |
| (config-router)# network 10.1.1.0/24 area 0.0.0.0 | Add network address |
| (config-router)# network 30.1.1.0/24 area 0.0.0.0 | Add network address |
| (config-router)# network 50.1.1.0/24 area 0.0.0.0 | Add network address |
| (config-router)# network 60.1.1.0/24 area 0.0.0.0 | Add network address |
| (config-router)# segment-routing mpls | Enable SR ospf |

P2:

| | |
|---|---------------------------------|
| (config)#segment-routing | Configure SR. |
| (config-sr)# mpls sr-prefer | Prefer SR over other protocol. |
| (config)#exit | Exit sr mode. |
| (config-sr)#router ldp | Configure router ldp |
| (config-router)# transport-address ipv4 4.4.4.4 | Configure transport address. |
| (config-router)#exit | Exit |
| (config)#interface ce1 | Configure network interface ce1 |
| (config-if)# ip address 20.1.1.2/24 | Configure ipv4 address |
| (config-if)# mtu 1522 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |
| (config-if)#interface lo | Configure loopback interface |
| (config-if)# ip address 4.4.4.4/32 secondary | Configure secondary ip |

| | |
|---|----------------------------------|
| (config-if)# prefix-sid absolute 16004 | Configure SR segment id |
| (config-if)#interface xe14 | Configure network interface xe14 |
| (config-if)# ip address 50.1.1.2/24 | Configure ipv4 address |
| (config-if)# mtu 1522 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |
| (config-if)#interface xe20 | Configure network interface xe20 |
| (config-if)# ip address 70.1.1.1/24 | Configure ipv4 address |
| (config-if)# mtu 1522 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |
| (config-if)#interface xe26 | Configure network interface xe26 |
| (config-if)# ip address 40.1.1.2/24 | Configure ipv4 address |
| (config-if)# mtu 2000 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |
| (config-if)#router ospf 1 | Configure router ospf 1 |
| (config-router)# network 4.4.4.4/32 area 0.0.0.0 | Add loopback ip to ospf |
| (config-router)# network 20.1.1.0/24 area 0.0.0.0 | Add network address |
| (config-router)# network 40.1.1.0/24 area 0.0.0.0 | Add network address |
| (config-router)# network 50.1.1.0/24 area 0.0.0.0 | Add network address |
| (config-router)# network 70.1.1.0/24 area 0.0.0.0 | Add network address |
| (config-router)# segment-routing mpls | Enable SR ospf |

PE3

| | |
|--|--|
| Loopback Interface: | |
| #configure terminal | Enter configuration mode. |
| (config-if)#interface lo | Enter the Interface mode for the loopback interface. |
| (config-if)# ip address 5.5.5.5/32 secondary | Configure IP address on loopback interface. |
| (config-if)# prefix-sid absolute 16005 | Configure SR segment id |
| (config)#commit | Commit the transaction. |

Global EVPN MPLS Command:

| | |
|---|--|
| (config)#evpn mpls enable | Enable EVPN MPLS |
| (config)#evpn mpls irb | Enable EVPN MPLS IRB |
| (config)#commit | Commit candidate configuration to be running configuration Note: Reload is required after Enabling/Disabling EVPN MPLS Feature |
| (config)#evpn mpls vtep-ip-global 5.5.5.5 | Configuring VTEP global IP to loopback IP |
| (config-evpn-mpls)#commit | Commit the transaction. |

IP VRF Configuration:

| | |
|---|---|
| (config-vrf)#ip vrf vrf102 | Enter VRF mode |
| (config-vrf)# rd 10001:102 | Configuring Route-Distinguisher value |
| (config-vrf)# route-target both 1.1.1.1:102 | Configure RT value |
| (config-vrf)# l3vni 102 | Configure L3VNI to send ipv4 route through evpn |
| (config-vrf)#ip vrf vrf101 | Create ip vrf vrf101 |
| (config-vrf)# rd 10002:1 | Configuring Route-Distinguisher value |
| (config-vrf)# route-target both 1.1.1.1:1 | Configure RT value |
| (config-vrf)# l3vni 101 | Configure L3VNI to send ipv4 route through evpn |
| (config)#commit | Commit the transaction. |

Configure SR:

| | |
|---|---|
| (config)#segment-routing | Configure SR |
| (config-sr)# mpls sr-prefer | Prefer SR over other protocols for mpls |
| (config)#commit | Commit the transaction. |
| Global LDP: | |
| (config)#router ldp | Enter the Router LDP mode. |
| (config-router)# transport-address ipv4 5.5.5.5 | Configure transport address |
| (config-router)#exit | Exit |
| (config)#commit | Commit the transaction. |

Interface Configuration Network Side:

| | |
|-------------------------------------|---------------------------------|
| (config-if)#interface ce3 | Configure network interface ce3 |
| (config-if)# ip address 60.1.1.2/24 | Configure ipv4 address |
| (config-if)# mtu 1522 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |

| | |
|---|------------------------------------|
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |
| (config-if)#interface xe20 | Enter the Interface mode for xe20. |
| (config-if)# ip address 70.1.1.2/24 | Configure ipv4 address |
| (config-if)# mtu 1522 | Configure MTU |
| (config-if)# label-switching | Configure label switching |
| (config-if)# mpls ldp-igp sync ospf | Configure ldp ospf sync |
| (config-if)# ip ospf network point-to-point | Configure ospf p2p |
| (config-if)# enable-ldp ipv4 | Enable ldp |
| (config)#commit | Commit the transaction. |

OSPF Configuration:

| | |
|---|---|
| (config)#router ospf 1 | Enter the Router OSPF mode. |
| (config-router)# network 5.5.5.5/32 area 0.0.0.0 | Advertise loopback address in OSPF. |
| (config-router)# network 60.1.1.0/24 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# network 70.1.1.0/24 area 0.0.0.0 | Advertise network address in OSPF. |
| (config-router)# segment-routing mpls | Enable SR ospf |
| (config-router)#exit | Exit Router OSPF mode and return to Configure mode. |
| (config)#commit | Commit the transaction. |

BGP Configuration:

| | |
|--|---|
| (config-router)#router bgp 100 | Enter the Router BGP mode, ASN: 100 |
| (config-router)# neighbor 1.1.1.1 remote-as 100 | Configuring PE1 as iBGP neighbor using it's loopback IP |
| (config-router)# neighbor 1.1.1.1 fall-over bfd multihop | Configure BFD |
| (config-router)# neighbor 2.2.2.2 remote-as 100 | Configuring PE2 as iBGP neighbor using it's loopback IP |
| (config-router)# neighbor 2.2.2.2 fall-over bfd multihop | Configure BFD |
| (config-router)# neighbor 1.1.1.1 update-source lo | Source of routing updates as loopback |
| (config-router)# neighbor 2.2.2.2 update-source lo | Source of routing updates as loopback |
| (config-router)# address-family ipv4 unicast | Enter ipv4 unicast address family |
| (config-router-af)# neighbor 1.1.1.1 activate | Activate neighbor PE1 |
| (config-router-af)# neighbor 2.2.2.2 activate | Activate neighbor PE2 |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family l2vpn evpn | Enter evpn address family |

| | |
|---|--|
| (config-router-af)# neighbor 1.1.1.1 activate | Activate neighbor PE1 |
| (config-router-af)# neighbor 2.2.2.2 activate | Activate neighbor PE2 |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv6 unicast | Enter ipv6 unicast address family |
| (config-router-af)# neighbor 1.1.1.1 activate | Activate neighbor PE1 |
| (config-router-af)# neighbor 2.2.2.2 activate | Activate neighbor PE2 |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv4 vrf vrf102 | Enter vrf vrf102 address family |
| (config-router-af)# redistribute connected | Redistribute vrf connected routes to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv6 vrf vrf102 | Enter ipv6 vrf vrf102 address family |
| (config-router-af)# redistribute connected | Redistribute vrf connected routes to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv4 vrf vrf101 | Enter vrf vrf102 address family |
| (config-router-af)# redistribute connected | Redistribute vrf connected routes to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)# address-family ipv6 vrf vrf101 | Enter ipv6 vrf vrf101 address family |
| (config-router-af)# redistribute connected | Redistribute vrf connected routes to bgp |
| (config-router-af)# exit-address-family | Exit address family |
| (config-router)#exit | Exit bgp mode |
| (config)#commit | Commit the transaction. |

L3 VRF Port Configuration:

| | |
|--|--|
| (config-if)#interface xe13 | Enter the Interface mode for xe3. This is for untagged traffic |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# ip address 201.1.1.1/24 | Assign ipv4 address |
| (config-if)# ipv6 address 2001::1/64 | Assign ipv6 address |
| (config-if)#interface xe13.1 | Create subinterface xe13.1 |
| (config-if)# encapsulation dot1q 10 | Configure encapsulation single tagged |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# ip address 210.1.1.1/24 | Assign ipv4 address |
| (config-if)# ipv6 address 2100::1/64 | Assign ipv6 address |
| (config-if)#interface xe13.2 | Create subinterface xe13.2 |
| (config-if)# encapsulation dot1q 11 inner-dot1q 11 | Configure encapsulation double tagged |
| (config-if)# ip vrf forwarding vrf101 | Map the vrf vrf101 |
| (config-if)# ip address 220.1.1.1/24 | Assign ipv4 address |

| | |
|--------------------------------------|-------------------------|
| (config-if)# ipv6 address 2200::1/64 | Assign ipv6 address |
| (config)#commit | Commit the transaction. |

Validation

PE1

```

7030-PE1#sh mlag domain summary
-----
Domain Configuration
-----

Domain System Number      : 1
Domain Address            : 1111.2222.3333
Domain Priority            : 32768
Intra Domain Interface    : xe4
Domain Adjacency          : UP
Domain Sync via           : Intra-domain-interface
-----

MLAG Configuration
-----

MLAG-1
Mapped Aggregator         : po1
Physical properties Digest : a2 58 27 76 9f 45 ff 6c 2a 62 65 aa b6 22 8f 81
Total Bandwidth           : 0
Mlag Sync                 : IN_SYNC
Mode                      : Active-Standby
Current Mlag state        : Active
Switchover-mode           : Revertive

7030-PE1#
Note:
PE1#sh etherchannel summary
  Aggregator po1 100001
  Aggregator Type: Layer2
  Admin Key: 16385 - Oper Key 16385
    Link: xe10 (5011) sync: 1 (Mlag-Active-link)
-----
  Aggregator po2 100002
  Aggregator Type: Layer3
  Admin Key: 0002 - Oper Key 0002
    Link: xe18 (10025) sync: 1
    Link: xe19 (10026) sync: 1
PE1#sh evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      evpn-id
=====
1.1.1.1     5.5.5.5       Installed   02:22:38     02:22:38     102

Total number of entries are 1
PE1#sh ip bgp vrf all
BGP table version is 1, local router ID is 120.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

      Network      Next Hop      Metric      LocPrf      Weight Path

BGP Route Table for VRF vrf101
*>i 45.1.1.0/24     5.5.5.5        0           100         0      ?
*>i 52.1.1.0/24     5.5.5.5        0           100         0      ?
*> 100.1.1.0/24     0.0.0.0        0           100         32768  ?
*>i 101.101.101.2/32 5.5.5.5        0           100         0      ?
*> 101.101.101.101/32

```

```

0.0.0.0          0          100      32768 ?
*> 110.1.1.0/24  0.0.0.0          0          100      32768 ?
*> 111.1.1.0/24  0.0.0.0          0          100      32768 ?
*                  0.0.0.0          1          100      32768 ?
*> 120.1.1.0/24  0.0.0.0          0          100      32768 ?
*>i 201.1.1.0     5.5.5.5          0          100      0        ?
*>i 210.1.1.0     5.5.5.5          0          100      0        ?
*>i 220.1.1.0     5.5.5.5          0          100      0        ?

```

Total number of prefixes 11

BGP Route Table for VRF vrf102

```

*>i 45.1.1.0/24  5.5.5.5          0          100      0        ?
*> 51.1.1.0/24  0.0.0.0          0          100      32768 ?
*>i 52.1.1.0/24  5.5.5.5          0          100      0        ?
*>i 101.101.101.2/32 5.5.5.5          0          100      0        ?
*>i 201.1.1.0     5.5.5.5          0          100      0        ?
*>i 210.1.1.0     5.5.5.5          0          100      0        ?
*>i 220.1.1.0     5.5.5.5          0          100      0        ?

```

Total number of prefixes 7

```

*>i 5.5.5.5/32    5.5.5.5          0          100      0        i

```

Total number of prefixes 1

PE1#sh bgp l2vpn evpn prefix-route

RD[10002:1]

| ESI | Eth-Tag | Prefix-Length | IP-Address | GW-IPAddress | L3VNID/LABEL | Nexthop | Encap | Router- |
|-------|---------|---------------|------------|--------------|--------------|---------|-------|-----------|
| Mac | | | | | | | | |
| 0 | 0 | 24 | 201.1.1.0 | 0.0.0.0 | 17 | 5.5.5.5 | MPLS | 0000:0000 |
| :0000 | | | | | | | | |
| 0 | 0 | 24 | 210.1.1.0 | 0.0.0.0 | 17 | 5.5.5.5 | MPLS | 0000:0000 |
| :0000 | | | | | | | | |
| 0 | 0 | 24 | 220.1.1.0 | 0.0.0.0 | 17 | 5.5.5.5 | MPLS | 0000:0000 |
| :0000 | | | | | | | | |
| 0 | 0 | 64 | 2001:: | :: | 17 | 5.5.5.5 | MPLS | 0000:0000 |
| :0000 | | | | | | | | |
| 0 | 0 | 64 | 2100:: | :: | 17 | 5.5.5.5 | MPLS | 0000:0000 |
| :0000 | | | | | | | | |
| 0 | 0 | 64 | 2200:: | :: | 17 | 5.5.5.5 | MPLS | 0000:0000 |
| :0000 | | | | | | | | |

PE1#sh mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM
 K - CLI ILM, T - MPLS-TP, s - Stitched ILM
 S - SNMP, L - LDP, R - RSVP, C - CRLDP
 B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
 O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
 P - SR Policy, U - unknown

LDP ilm-ecmp - disabled

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|-------------|----------|-----------|---------|--------------|
| Intf/VRF | Nexthop | | pri | LSP-Type | | |
| O> | 5.5.5.5/32 | 7 | 16005 | 16005 | N/A | po2 10.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| O> | 3.3.3.3/32 | 5 | 16003 | 3 | N/A | po2 10.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| B> | vrf101 | 2 | 17 | Nolabel | N/A | vrf101 N/A |
| | Yes | LSP_DEFAULT | | | | |
| B> | vrf102 | 1 | 16 | Nolabel | N/A | vrf102 N/A |
| | Yes | LSP_DEFAULT | | | | |
| O> | 4.4.4.4/32 | 6 | 16004 | 3 | N/A | ce0 20.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| L> | 4.4.4.4/32 | 10 | 35202 | 3 | N/A | ce0 20.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| L> | 3.3.3.3/32 | 8 | 35200 | 3 | N/A | po2 10.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| L> | 60.1.1.0/24 | 9 | 35201 | 3 | N/A | po2 10.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| O> | 20.1.1.2/32 | 3 | 35840 | 3 | N/A | ce0 20.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |

```

L> 70.1.1.0/24      12      35204      3      N/A      ce0      20.1.1.2
    Yes      LSP_DEFAULT
O> 10.1.1.2/32      4      35841      3      N/A      po2      10.1.1.2
    Yes      LSP_DEFAULT
PE1#sh mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
        B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
        (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

| Code | FEC | Nexthop | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|-------------|---------|--------|----------|-----------|-----|----------|-----------|------|
| O> | 3.3.3.3/32 | | 1 | 5 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | po2 | No | 10.1.1.2 | | | | | |
| L | 3.3.3.3/32 | | 9 | 4 | - | Yes | LSP_ | | |
| DEFAULT | 3 | po2 | No | 10.1.1.2 | | | | | |
| O> | 4.4.4.4/32 | | 2 | 6 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | ce0 | No | 20.1.1.2 | | | | | |
| L | 4.4.4.4/32 | | 4 | 3 | - | Yes | LSP_ | | |
| DEFAULT | 3 | ce0 | No | 20.1.1.2 | | | | | |
| O> | 5.5.5.5/32 | | 3 | 8 | 0 | Yes | LSP_ | | |
| DEFAULT | 16005 | po2 | No | 10.1.1.2 | | | | | |
| L | 5.5.5.5/32 | | 5 | 13 | - | Yes | LSP_ | | |
| DEFAULT | 34567 | po2 | No | 10.1.1.2 | | | | | |
| | | | | 14 | - | Yes | LSP_ | | |
| DEFAULT | 34566 | ce0 | No | 20.1.1.2 | | | | | |
| L> | 50.1.1.0/24 | | 7 | 15(e) | | | | | |
| | | | | 4 | - | Yes | LSP_ | | |
| DEFAULT | 3 | po2 | No | 10.1.1.2 | | | | | |
| | | | | 3 | - | Yes | LSP_ | | |
| DEFAULT | 3 | ce0 | No | 20.1.1.2 | | | | | |
| L> | 60.1.1.0/24 | | 10 | 16 | | | | | |
| | | | | 4 | - | Yes | LSP_ | | |
| DEFAULT | 3 | po2 | No | 10.1.1.2 | | | | | |
| L> | 70.1.1.0/24 | | 8 | 10 | | | | | |
| | | | | 3 | - | Yes | LSP_ | | |
| DEFAULT | 3 | ce0 | No | 20.1.1.2 | | | | | |

```

PE1#

```

PE2

```

PE2#show mlag domain summary

-----
Domain Configuration
-----

Domain System Number      : 2
Domain Address             : 1111.2222.3333
Domain Priority            : 32768
Intra Domain Interface    : xe4
Domain Adjacency          : UP
Domain Sync via           : Intra-domain-interface
-----

MLAG Configuration
-----

MLAG-1
Mapped Aggregator         : po1
Physical properties Digest : a2 58 27 76 9f 45 ff 6c 2a 62 65 aa b6 22 8f 81
Total Bandwidth           : 0
Mlag Sync                 : IN_SYNC
Mode                      : Active-Standby
Current Mlag state        : Standby
Switchover-mode           : Revertive

PE2#show evpn mpls tunnel

```


EVPN-MPLS Network tunnel Entries

| Source | Destination | Status | Up/Down | Update | evpn-id |
|---------|-------------|-----------|----------|----------|---------|
| 2.2.2.2 | 5.5.5.5 | Installed | 1d02h23m | 1d02h23m | 101 |

Total number of entries are 1

PE2#show evpn mpls tunnel label

EVPN-MPLS Network tunnel labels

(*) in Policy - tunnel-policy inherited from mac-vrf

| | | | | Local | Remote | MP |
|--------------|-----------|---------|----------|----------|----------|------|
| LS-Multipath | Underlay | | | | | |
| Destination | Status | VPN-ID | Policy | MC-Label | UC-Label | Grp- |
| Name | NHLFE-ix | NW-Intf | NW-Label | | | |
| 5.5.5.5 | Installed | 101 | -- | -- | 17 | -- |

Total number of entries are 1

PE2#show bgp l2vpn evpn prefix-route

RD[10002:1]

| ESI | Eth-Tag | Prefix-Length | IP-Address | GW- | | | | |
|-------------|--------------|---------------|---------------|------------|----|---------|------|-----|
| IPAddress | L3VNID/LABEL | NextHop | Encap | Router-Mac | | | | |
| 0 | 0 | 24 | 45.1.1.0 | 0.0.0.0 | 17 | 5.5.5.5 | MPLS | 000 |
| 0:0000:0000 | | | | | | | | |
| 0 | 0 | 24 | 52.1.1.0 | 0.0.0.0 | 17 | 5.5.5.5 | MPLS | 000 |
| 0:0000:0000 | | | | | | | | |
| 0 | 0 | 24 | 201.1.1.0 | 0.0.0.0 | 17 | 5.5.5.5 | MPLS | 000 |
| 0:0000:0000 | | | | | | | | |
| 0 | 0 | 24 | 210.1.1.0 | 0.0.0.0 | 17 | 5.5.5.5 | MPLS | 000 |
| 0:0000:0000 | | | | | | | | |
| 0 | 0 | 24 | 220.1.1.0 | 0.0.0.0 | 17 | 5.5.5.5 | MPLS | 000 |
| 0:0000:0000 | | | | | | | | |
| 0 | 0 | 32 | 101.101.101.2 | 0.0.0.0 | 17 | 5.5.5.5 | MPLS | 000 |
| 0:0000:0000 | | | | | | | | |
| 0 | 0 | 64 | 2001:: | :: | 17 | 5.5.5.5 | MPLS | 000 |
| 0:0000:0000 | | | | | | | | |
| 0 | 0 | 64 | 2100:: | :: | 17 | 5.5.5.5 | MPLS | 000 |
| 0:0000:0000 | | | | | | | | |
| 0 | 0 | 64 | 2200:: | :: | 17 | 5.5.5.5 | MPLS | 000 |
| 0:0000:0000 | | | | | | | | |

PE2#show ip bgp vrf all

BGP table version is 1, local router ID is 0.0.0.0

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
l - labeled, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|--------------------------------|----------|--------|--------|--------|------|
| BGP Route Table for VRF vrf101 | | | | | |
| *>i 201.1.1.0 | 5.5.5.5 | 0 | 100 | 0 | ? |
| *>i 210.1.1.0 | 5.5.5.5 | 0 | 100 | 0 | ? |
| *>i 220.1.1.0 | 5.5.5.5 | 0 | 100 | 0 | ? |

Total number of prefixes 6

| | | | | | |
|----------------|---------|---|-----|---|---|
| *>i 5.5.5.5/32 | 5.5.5.5 | 0 | 100 | 0 | i |
|----------------|---------|---|-----|---|---|

Total number of prefixes 1

PE2#show mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM

K - CLI ILM, T - MPLS-TP, s - Stitched ILM

S - SNMP, L - LDP, R - RSVP, C - CRLDP

B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT

O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI

P - SR Policy, U - unknown

LDP ilm-ecmp - disabled

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|-------------|----------|-----------|---------|---------------|
| Intf/VRF | Nexthop | | pri | LSP-Type | | |
| O> | 5.5.5.5/32 | 8 | 16005 | 16005 | N/A | xe6 30.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| O> | 1.1.1.1/32 | 5 | 16001 | 16001 | N/A | xe6 30.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| B> | vrf101 | 2 | 17 | Nolabel | N/A | vrf101 N/A |
| | Yes | LSP_DEFAULT | | | | |
| B> | vrf102 | 1 | 16 | Nolabel | N/A | vrf102 N/A |
| | Yes | LSP_DEFAULT | | | | |
| O> | 4.4.4.4/32 | 7 | 16004 | 3 | N/A | xe26 40.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| O> | 3.3.3.3/32 | 6 | 16003 | 3 | N/A | xe6 30.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| L> | 10.1.1.0/24 | 13 | 35204 | 3 | N/A | xe6 30.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| L> | 20.1.1.0/24 | 10 | 35201 | 3 | N/A | xe26 40.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| L> | 4.4.4.4/32 | 9 | 35200 | 3 | N/A | xe26 40.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| L> | 3.3.3.3/32 | 12 | 35203 | 3 | N/A | xe6 30.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| L> | 70.1.1.0/24 | 11 | 35202 | 3 | N/A | xe26 40.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| O> | 30.1.1.2/32 | 3 | 35840 | 3 | N/A | xe6 30.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| L> | 60.1.1.0/24 | 14 | 35205 | 3 | N/A | xe6 30.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |
| O> | 40.1.1.2/32 | 4 | 35841 | 3 | N/A | xe26 40.1.1.2 |
| | Yes | LSP_DEFAULT | | | | |

PE2#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
 (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

| Code | FEC | FTN-ID | Nhife-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|-------------|--------|----------|-----------|-----|----------|-----------|------|
| Intf | ELC Nexthop | | | | | | | |
| O> | 1.1.1.1/32 | 1 | 6 | 0 | Yes | LSP_ | | |
| DEFAULT | 16001 xe6 | No | 30.1.1.2 | | | | | |
| L | 1.1.1.1/32 | 5 | 16 | - | Yes | LSP_ | | |
| DEFAULT | 34573 xe26 | No | 40.1.1.2 | | | | | |
| | | | 17 | - | Yes | LSP_ | | |
| DEFAULT | 34573 xe6 | No | 30.1.1.2 | | | | | |
| O> | 3.3.3.3/32 | 2 | 7 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 xe6 | No | 30.1.1.2 | | | | | |
| L | 3.3.3.3/32 | 6 | 3 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe6 | No | 30.1.1.2 | | | | | |
| O> | 4.4.4.4/32 | 3 | 8 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 xe26 | No | 40.1.1.2 | | | | | |
| L | 4.4.4.4/32 | 11 | 4 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe26 | No | 40.1.1.2 | | | | | |
| O> | 5.5.5.5/32 | 4 | 10 | 0 | Yes | LSP_ | | |
| DEFAULT | 16005 xe6 | No | 30.1.1.2 | | | | | |
| L | 5.5.5.5/32 | 7 | 19 | - | Yes | LSP_ | | |
| DEFAULT | 34575 xe26 | No | 40.1.1.2 | | | | | |
| | | | 20 | - | Yes | LSP_ | | |
| DEFAULT | 34575 xe6 | No | 30.1.1.2 | | | | | |
| L> | 10.1.1.0/24 | 8 | 14 | | | | | |
| | | | 3 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe6 | No | 30.1.1.2 | | | | | |
| L> | 20.1.1.0/24 | 12 | 21 | | | | | |
| | | | 4 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe26 | No | 40.1.1.2 | | | | | |
| L> | 50.1.1.0/24 | 9 | 22 (e) | | | | | |
| | | | 4 | - | Yes | LSP_ | | |
| DEFAULT | 3 xe26 | No | 40.1.1.2 | | | | | |

```

3 - Yes LSP_
DEFAULT 3 xe6 No 30.1.1.2
L> 60.1.1.0/24 10 14
3 - Yes LSP_
DEFAULT 3 xe6 No 30.1.1.2
L> 70.1.1.0/24 13 21
4 - Yes LSP_
DEFAULT 3 xe26 No 40.1.1.2
PE2#

PE3:

PE3#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source Destination Status Up/Down Update evpn-id
=====
5.5.5.5 1.1.1.1 Installed 2d05h17m 2d05h17m 102
5.5.5.5 1.1.1.1 Installed 2d05h17m 2d05h17m 101

Total number of entries are 2
PE3#show ip bgp vrf all
BGP table version is 1, local router ID is 220.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path
BGP Route Table for VRF vrf101
*> 45.1.1.0/24 210.1.1.2 0 100 32768 ?
*>i 51.1.1.0/24 1.1.1.1 0 100 0 ?
*> 52.1.1.0/24 0.0.0.0 0 100 32768 ?
*>i 100.1.1.0/24 1.1.1.1 0 100 0 ?
*> 101.101.101.2/32 0.0.0.0 0 100 32768 ?
*>i 101.101.101.101/32
1.1.1.1 0 100 0 ?
*>i 110.1.1.0/24 1.1.1.1 0 100 0 ?
*>i 111.1.1.0/24 1.1.1.1 0 100 0 ?
*>i 120.1.1.0/24 1.1.1.1 0 100 0 ?
*> 201.1.1.0 0.0.0.0 0 100 32768 ?
*> 210.1.1.0 0.0.0.0 0 100 32768 ?
*> 220.1.1.0 0.0.0.0 0 100 32768 ?

Total number of prefixes 12

...skipping 1 line
BGP Route Table for VRF vrf102
*>i 51.1.1.0/24 1.1.1.1 0 100 0 ?

Total number of prefixes 1
*> 5.5.5.5/32 0.0.0.0 0 100 32768 i

Total number of prefixes 1

PE3#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown

LDP ilm-ecmp - disabled
Code FEC/VRF/L2CKT ILM-ID In-Label Out-Label In-Intf Out-
Intf/VRF Nexthop pri LSP-Type
O> 4.4.4.4/32 6 16004 3 N/A xe20 70.1.1.1
Yes LSP_DEFAULT
B> vrf101 2 17 Nolabel N/A vrf101 N/A
Yes LSP_DEFAULT

```

```

B> vrf102 1 16 Nolabel N/A vrf102 N/A
Yes LSP_DEFAULT
O> 3.3.3.3/32 5 16003 3 N/A ce3 60.1.1.1
Yes LSP_DEFAULT
O> 1.1.1.1/32 12 16001 16001 N/A ce3 60.1.1.1
Yes LSP_DEFAULT
L> 20.1.1.0/24 10 35203 3 N/A xe20 70.1.1.1
Yes LSP_DEFAULT
L> 4.4.4.4/32 8 35201 3 N/A xe20 70.1.1.1
Yes LSP_DEFAULT
L> 3.3.3.3/32 7 35200 3 N/A ce3 60.1.1.1
Yes LSP_DEFAULT
L> 40.1.1.0/24 9 35202 3 N/A xe20 70.1.1.1
Yes LSP_DEFAULT
O> 70.1.1.1/32 4 35840 3 N/A xe20 70.1.1.1
Yes LSP_DEFAULT
L> 10.1.1.0/24 11 35204 3 N/A ce3 60.1.1.1
Yes LSP_DEFAULT
O> 60.1.1.1/32 3 35841 3 N/A ce3 60.1.1.1
Yes LSP_DEFAULT
PE3#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|-------------|---------|----------|-----------|-----|----------|-----------|------|
| Intf | ELC | Nexthop | | | | | | |
| O> | 1.1.1.1/32 | 9 | 13 | 0 | Yes | LSP_ | | |
| DEFAULT | 16001 | ce3 | No | 60.1.1.1 | | | | |
| L | 1.1.1.1/32 | 10 | 14 | - | Yes | LSP_ | | |
| DEFAULT | 34566 | ce3 | No | 60.1.1.1 | | | | |
| | | | | 16 | Yes | LSP_ | | |
| DEFAULT | 34564 | xe20 | No | 70.1.1.1 | | | | |
| O> | 3.3.3.3/32 | 1 | 6 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | ce3 | No | 60.1.1.1 | | | | |
| L | 3.3.3.3/32 | 3 | 3 | - | Yes | LSP_ | | |
| DEFAULT | 3 | ce3 | No | 60.1.1.1 | | | | |
| O> | 4.4.4.4/32 | 2 | 7 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe20 | No | 70.1.1.1 | | | | |
| L | 4.4.4.4/32 | 5 | 4 | - | Yes | LSP_ | | |
| DEFAULT | 3 | xe20 | No | 70.1.1.1 | | | | |
| L> | 10.1.1.0/24 | 8 | 8 | - | Yes | LSP_ | | |
| | | | | 3 | | | | |
| DEFAULT | 3 | ce3 | No | 60.1.1.1 | | | | |
| L> | 20.1.1.0/24 | 7 | 10 | - | Yes | LSP_ | | |
| | | | | 4 | | | | |
| DEFAULT | 3 | xe20 | No | 70.1.1.1 | | | | |
| L> | 40.1.1.0/24 | 6 | 10 | - | Yes | LSP_ | | |
| | | | | 4 | | | | |
| DEFAULT | 3 | xe20 | No | 70.1.1.1 | | | | |
| L> | 50.1.1.0/24 | 4 | 11(e) | - | Yes | LSP_ | | |
| | | | | 4 | | | | |
| DEFAULT | 3 | xe20 | No | 70.1.1.1 | | | | |
| | | | | 3 | Yes | LSP_ | | |
| DEFAULT | 3 | ce3 | No | 60.1.1.1 | | | | |

PE3

```
PE3#show bgp l2vpn evpn prefix-route
```

| ESI | Eth-Tag | Prefix-Length | IP-Address | GW- |
|-----------|----------------|---------------|------------|--------------------|
| IPAddress | L3VNID/LABEL | Nexthop | Router-Mac | |
| 0 | 0 | 24 | 100.1.1.0 | 0.0.0.0 17 1.1.1.1 |
| MPLS | 0000:0000:0000 | | | |
| 0 | 0 | 24 | 110.1.1.0 | 0.0.0.0 17 1.1.1.1 |
| MPLS | 0000:0000:0000 | | | |

```

0          0          24          111.1.1.0          0.0.0.0          17          1.1.1.1
MPLS      0000:0000:0000
0          0          24          120.1.1.0          0.0.0.0          17          1.1.1.1
MPLS      0000:0000:0000
0          0          64          1000::          ::          17          1.1.1.1
MPLS      0000:0000:0000
0          0          64          1100::          ::          17          1.1.1.1
MPLS      0000:0000:0000
0          0          64          1110::          ::          17          1.1.1.1
MPLS      0000:0000:0000
0          0          64          1200::          ::          17          1.1.1.1
MPLS      0000:0000:0000

RD[10001:102]
ESI          Eth-Tag Prefix-Length IP-Address  GW-
IPAddress  L3VNID/LABEL  Nexthop  Encap  Router-Mac
0          0          24          51.1.1.0          0.0.0.0          16          1.1.1.1
MPLS      0000:0000:0000
PE3#

```

L3VPN GR Configuration

Using BGP graceful restart, the data-forwarding plane of a router can continue to process and forward packets even if the control plane - which is responsible for determining best paths - fails. Graceful restart also reduces routing flaps, stabilizing the network and reducing control-plane resource consumption.

By exchanging a new BGP capability (BGP capability code 64) in the initial BGP open messages that establish the session, the restarting router and its peers show that they are aware of the BGP graceful restart mechanism when the initial BGP connection is established. In addition, the restarting router provides its peers with a list of supported address-families (VPNv4, IPv4, and IPv6) for which it can maintain a forwarding state across a BGP restart.

The peer router's TCP connection might be cleared, when the router's BGP process is restarted. Under normal circumstances, this would cause the peer router to clear all routes associated with the restarting router. But with a BGP graceful restart, this doesn't happen. Instead, in expectation of the restarting router shortly re-establishing the BGP session, the peer router marks all routes as "stale" yet continues to use them to forward packets. Likewise, the restarting router also continues forwarding packets in the interim.

When the restarting router opens the new BGP session, it will again send BGP capability 64 to its peers. But this time, flags will be set in the graceful restart capabilities exchange to let the peer router know that the BGP process has restarted.

The goal of the BGP graceful restart was to minimize the duration and reach of an outage associated with a failed BGP process. To do this, the software extensions must be deployed on both the router restarting the BGP process and the BGP peers of that router. The peers help the BGP process regain lost forwarding information and also help isolate failures from the rest of the network.


While forwarding packets, the peer router will refresh the restarting router with any relevant BGP routing information base (RIB) updates. The peer signals that it has finished sending the updates with an "End-of-RIB" (EOR) marker - an "empty" BGP update message. EOR markers help speed convergence because once the restarting router has received them from all peers, it can begin best-path selection again using the new routing information. Similarly, the restarting router then sends any updates to its peer routers and uses the EOR marker to indicate the completion of the process.

As part of this feature, we will be extending the feature for VPNv4 AF.

Topology

In the below example shows to configure BGP VPNv4 neighborhood between PE1 and PE2.

Figure 98. L3VPN GR Topology



Configuration

Below are the configurations and validations of L3VPN GR with OSPF as IGP. We can also configure ISIS as IGP and LDP/RSVP as transport.


PE1

| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 1.1.1.1/32 secondary | Set the IP address of the loopback interface to 1.1.1.1/32 |
| (config-if)#exit | Exit interface mode. |
| (config)#ip-vrf l3vpn | Configure IP VRF L3VPN. |
| (config-vrf)#rd 1:300 | Enter RD value. |
| (config-vrf)#route-target both 300:400 | Enter RT value. |
| (config-vrf)#exit | Exiting from VRF mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 1.1.1.1 | Set the router ID to IP address 1.1.1.1. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Configure targeted peer. |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit-targeted-peer-mode |
| (config-router)#targeted-peer ipv4 5.5.5.5 | Configure targeted peer. |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ip address 10.10.10.1/24 | Configure IPv4 address for xe1. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe1. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe2.1 | Configure access-port . |
| (config-if)#description MPLS-L3VPN | Giving Interface Description |
| (config-if)#encapsulation dot1q 4 | Setting Encapsulation to dot1q with VLAN ID 4 |
| (config-if)#load-interval 30 | Load interval setting |
| (config-if)#ip vrf forwarding l3vpn | Bind the interface connected to the CE router with VRF l3vpn |
| (config-if)#ip address 172.168.25.2/24 | Assign the IPv4 address. |
| (config)#router ospf 1 | Configure the routing process and specify the Process ID 100. The Process ID should be a unique positive integer identifying the routing process. |

| | |
|---|---|
| (config)#ospf router-id 1.1.1.1 | Configure OSPF router-ID same as loopback interface IP address |
| (config-router) #network 1.1.1.1/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router) #network 10.10.10.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 100 | Enter router BGP mode |
| (config-router)#bgp router-id 1.1.1.1 | Configuring the BGP router id 1.1.1.1. |
| (config-router)#bgp graceful-restart restart-time 100 | Enable BGP GR with restart timer 100. |
| (config-router)#neighbor 5.5.5.5 remote-as 100 | Configure neighbor 5.5.5.5. |
| (config-router)#neighbor 5.5.5.5 update-source lo | Update source lo for neighbor 5.5.5.5. |
| (config-router)#address-family ipv4 unicast | Enter address-family IPv4 unicast. |
| (config-router-af)#redistribute connected | Redistribute connected. |
| (config-router)#neighbor 5.5.5.5 activate | Activate neighbor. |
| (config-router)#address-family vpnv4 unicast | Entering Address family VPNv4 unicast. |
| (config-router-af)#neighbor 5.5.5.5 activate | Activate the neighbor 5.5.5.5. |
| (config-router-af)#neighbor 5.5.5.5 capability graceful-restart | Activate capability graceful restart for neighbor 5.5.5.5. |
| (config-router-af)#exit-address-family | Exit address family. |
| (config-router)#address-family ipv4 vrf l3vpn | Entering address family. |
| (config-router-af)#redistribute connected | Redistribute connected. |
| (config-router)#neighbor 172.168.25.1 remote-as 600 | Configure neighbor 172.168.25.1. |
| (config-router)#neighbor 172.168.25.1 activate | Activate neighbor. |
| (config-router-af)#commit | Commit all the transactions. |

P1

| | |
|--|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 2.2.2.2/32 secondary | Set the IP address of the loopback interface to 2.2.2.2/32. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2. |
| (config-router)#transport-address ipv4 2.2.2.2 0 | Configure the transport address for IPV4 (for IPV6, use ipv6) to be used for a TCP session over which LDP will run. |

| | |
|---|---|
| |  Note: It is preferable to use the loopback address as the transport address. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit-targeted-peer-mode. |
| (config-router)#targeted-peer ipv4 5.5.5.5 | Configure targeted peer. |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit-targeted-peer-mode. |
| (config-router)#exit | Exit-targeted-peer-mode. |
| (config-if)#exit | Exit router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ip address 10.10.10.2/24 | Configure IPv4 address for xe1. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe1. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip address 40.40.40.1/24 | Configure IPv4 address for xe2. |
| (config-if)#label-switching | Enable label switching on interface xe2. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe2. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 1 | Configure the routing process and specify the Process ID. The Process ID should be a unique positive integer identifying the routing process. |
| (config)#ospf router-id 2.2.2.2 | Configure OSPF router-ID same as loopback interface IP address. |
| (config-router) #network 2.2.2.2/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router) #network 10.10.10.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#network 40.40.40.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router)#bfd all-interfaces | Enable the OSPF enabled interfaces with BFD. |
| (config-if)#exit | Exit interface mode. |

PE2

| | |
|---|---|
| #configure terminal | Enter configuration mode. |
| (config)#interface lo | Specify the loopback (lo) interface to be configured. |
| (config-if)#ip address 5.5.5.5/32 secondary | Set the IP address of the loopback interface to 5.5.5.5/32. |

| | |
|---|---|
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router mode for LDP. |
| (config-router)#router-id 5.5.5.5 | Set the router ID to IP address 5.5.5.5. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit-targeted-peer-mode. |
| (config-router)#targeted-peer ipv4 2.2.2.2 | Configure targeted peer. |
| (config-router-targeted-peer)#exit-targeted-peer-mode | Exit-targeted-peer-mode. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ip address 40.40.40.2/24 | Configure IPv4 address for xe1. |
| (config-if)#label-switching | Enable label switching on interface xe1. |
| (config-if)#enable-ldp ipv4 | Enable LDP for IPv4 on xe1. |
| (config-if)#exit | Exit interface mode. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2.1 | Enter interface mode. |
| (config-if)#description MPLS-L3VPN | Giving Interface Description. |
| (config-if)#encapsulation dot1q 4 | Setting Encapsulation to dot1q with VLAN ID 4. |
| (config-if)#load-interval 30 | Load interval setting. |
| (config-if)#ip vrf forwarding l3vpn | Bind the interface connected to the CE router with VRF L3VPN. |
| (config-if)#ip address 172.168.25.2/24 | Assign the IPv4 address. |
| (config)#router ospf 1 | The Process ID should be a unique positive integer identifying the routing process. |
| (config)#ospf router-id 5.5.5.5 | Configure OSPF router-ID same as loopback interface IP address. |
| (config-router) #network 5.5.5.5/32 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-router) #network 40.40.40.0/24 area 0 | Define the interface on which OSPF runs and associate the area ID (0) with the interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 100 | Enter router BGP mode |
| (config-router)#bgp router-id 5.5.5.5 | Configuring the bgp router id 1.1.1.1. |
| (config-router)#bgp graceful-restart restart-time 100 | Enable BGP GR with restart timer 100. |
| (config-router)#neighbor 1.1.1.1 remote-as 100 | Configure neighbor 1.1.1.1. |
| (config-router)#neighbor 1.1.1.1 update-source lo | Update source lo for neighbor 1.1.1.1. |
| (config-router)#address-family ipv4 unicast | Enter address-family ipv4 unicast. |
| (config-router-af)#redistribute connected | Redistribute connected. |

| | |
|---|--|
| (config-router)#neighbor 1.1.1.1 activate | Activate neighbor. |
| (config-router)#address-family vpnv4 unicast | Entering Address family VPNv4 unicast. |
| (config-router-af)#neighbor 1.1.1.1 activate | Activate the neighbor 1.1.1.1. |
| (config-router-af)#neighbor 1.1.1.1 capability graceful-restart | Activate capability graceful restart for neighbor 1.1.1.1. |
| (config-router-af)#exit-address-family | Exit address family. |
| (config-router)#address-family ipv4 vrf l3vpn | Entering address family. |
| (config-router-af)#redistribute connected | Redistribute connected. |
| (config-router)#neighbor 172.168.26.1 remote-as 700 | Configure neighbor 172.168.26.1. |
| (config-router)#neighbor 172.168.26.1 activate | Activate neighbor. |
| (config-router-af)#commit | Commit all the transactions. |

Validation

Restart BGP Gracefully

PE1

```

PE1#restart bgp graceful
%Warning : BGP process will stop and needs to restart manually,
You may lose bgp configuration,if not saved
Proceed for graceful restart? (y/n):y
%% Managed module is down or crashed

R1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
      K - CLI ILM, T - MPLS-TP, s - Stitched ILM
      S - SNMP, L - LDP, R - RSVP, C - CRLDP
      B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
      O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
      P - SR Policy, U - unknown

Code   FEC/VRF/L2CKT   ILM-ID   In-Label   Out-Label   In-Intf   Out-
Intf/VRF   Nexthop       LSP-Type
      LSP_DEFAULT
B> p
77.77.80.0/24   7           24323     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B> p
77.77.78.0/24   5           24321     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B> p
77.77.77.0/24   4           24320     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B> p
77.77.79.0/24   6           24322     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B> p
77.77.81.0/24   8           24324     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B> p
172.168.25.0/24  9           24325     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT

V      12ckt:900       1           24960     Nolabel     po1       xe1       N/A
      LSP_DEFAULT

```

```

PE1#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, B - BGP FTN
(m) - Service mapped over multipath transport

Code      FEC                FTN-ID   Tunnel-id   Pri   LSP-Type   Out-Label   Out-
Intf      Nexthop

  B> p 88.88.88.0/24      1         0           Yes   LSP_DEFAULT 24321       -           5.5.5.5
  B>p 88.88.89.0/24      2         0          Yes   LSP_DEFAULT 24321       -           5.5.5.5
    B> p 88.88.90.0/24    3         0           Yes   LSP_DEFAULT 24321       -
    5.5.5.5
    B >p 88.88.91.0/24    4         0           Yes   LSP_DEFAULT 24321       -
    5.5.5.5
    B >p 88.88.92.0/24    5         0           Yes   LSP_DEFAULT 24321       -
    5.5.5.5
  B> p 172.168.26.0/24    6         0           Yes   LSP_DEFAULT 24321       -
    5.5.5.5

PE1#show nsm forwarding-timer
Protocol-Name GR-State Time Remaining (sec)   Disconnected-time
BGP           ACTIVE      74                2022/01/13 16:33:43

PE#show run bgp
!
PE1#show ip bgp vpnv4 all

```

PE2

```

PE2#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, l -
labeled

          S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric   LocPrf   Weight Path
Route Distinguisher: 1:300 (Default for VRF l3vpn)
*>i 77.77.77.0/24    1.1.1.1            0         100      0 600 i
*>i 77.77.78.0/24    1.1.1.1            0         100      0 600 i
*>i 77.77.79.0/24    1.1.1.1            0         100      0 600 i
*>i 77.77.80.0/24    1.1.1.1            0         100      0 600 i
*>i 77.77.81.0/24    1.1.1.1            0         100      0 600 i
*> 1 88.88.88.0/24    172.168.26.1       0         100      0 700 i
*> 1 88.88.89.0/24    172.168.26.1       0         100      0 700 i
*> 1 88.88.90.0/24    172.168.26.1       0         100      0 700 i
*> 1 88.88.91.0/24    172.168.26.1       0         100      0 700 i
*> 1 88.88.92.0/24    172.168.26.1       0         100      0 700 i
*>i 172.168.25.0/24  1.1.1.1            0         100      0 ?
*> 1 172.168.26.0/24 0.0.0.0            0         100     32768 ?
  Announced routes count = 6
  Accepted routes count = 6
Route Distinguisher: 1:300
S>i 77.77.77.0/24    1.1.1.1            0         100      0 600 i
S>i 77.77.78.0/24    1.1.1.1            0         100      0 600 i
S>i 77.77.79.0/24    1.1.1.1            0         100      0 600 i
S>i 77.77.80.0/24    1.1.1.1            0         100      0 600 i
S>i 77.77.81.0/24    1.1.1.1            0         100      0 600 i
S>i 172.168.25.0/24  1.1.1.1            0         100      0 ?
  Announced routes count = 0

```

After Restarting the BGP Manually

PE1

```

PE1#start-shell
bash-5.0$ su

```

```

Password:
root@PE1:/home/ocnos# cd /usr/local/sbin/
root@PE1:/usr/local/sbin# ./bgpd -d

PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
      K - CLI ILM, T - MPLS-TP, s - Stitched ILM
      S - SNMP, L - LDP, R - RSVP, C - CRLDP
      B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
      O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
      P - SR Policy, U - unknown

Code   FEC/VRF/L2CKT   ILM-ID   In-Label   Out-Label   In-Intf   Out-
Intf/VRF   Nexthop   LSP-Type
-----
B>
77.77.80.0/24   7         24323     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B>
77.77.78.0/24   5         24321     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B>
77.77.77.0/24   4         24320     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B>
77.77.79.0/24   6         24322     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B>
77.77.81.0/24   8         24324     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
B>
172.168.25.0/24  9         24325     Nolabel     N/A         13vpn     N/A
      LSP_DEFAULT
V   12ckt:900       1         24960     Nolabel     po1       xe1       N/A
      LSP_DEFAULT

PE1#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, B - BGP FTN
(m) - Service mapped over multipath transport

Code   FEC           FTN-ID   Tunnel-id   Pri   LSP-Type   Out-Label   Out-
Intf   Nexthop
-----
B>88.88.88.0/24   1         0         Yes   LSP_DEFAULT 24321     -           5.5.5.5
B>88.88.89.0/24   2         0         Yes   LSP_DEFAULT 24321     -           5.5.5.5
B>88.88.90.0/24   3         0         Yes   LSP_DEFAULT 24321     -           5.5.5.5
B>88.88.91.0/24   4         0         Yes   LSP_DEFAULT 24321     -           5.5.5.5
B>88.88.92.0/24   5         0         Yes   LSP_DEFAULT 24321     -           5.5.5.5
B> 172.168.26.0/24 6         0         Yes   LSP_DEFAULT 24321     -           5.5.5.5

```

PE2

```

PE2#show ip bgp vpnv4 all
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, l -
labeled
          S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Network      Next Hop      Metric   LocPrf   Weight Path
Route Distinguisher: 1:300 (Default for VRF 13vpn)
*>i 77.77.77.0/24  1.1.1.1       0       100      0 600 i
*>i 77.77.78.0/24  1.1.1.1       0       100      0 600 i
*>i 77.77.79.0/24  1.1.1.1       0       100      0 600 i
*>i 77.77.80.0/24  1.1.1.1       0       100      0 600 i
*>i 77.77.81.0/24  1.1.1.1       0       100      0 600 i
*> 1 88.88.88.0/24  172.168.26.1  0       100      0 700 i
*> 1 88.88.89.0/24  172.168.26.1  0       100      0 700 i
*> 1 88.88.90.0/24  172.168.26.1  0       100      0 700 i

```

```
*> 1 88.88.91.0/24 172.168.26.1 0 100 0 700 i
*> 1 88.88.92.0/24 172.168.26.1 0 100 0 700 i
*>i 172.168.25.0/24 1.1.1.1 0 100 0 ?
*> 1 172.168.26.0/24 0.0.0.0 0 100 32768 ?
  Announced routes count = 6
  Accepted routes count = 6
Route Distinguisher: 1:300
>i 77.77.77.0/24 1.1.1.1 0 100 0 600 i
>i 77.77.78.0/24 1.1.1.1 0 100 0 600 i
>i 77.77.79.0/24 1.1.1.1 0 100 0 600 i
>i 77.77.80.0/24 1.1.1.1 0 100 0 600 i
>i 77.77.81.0/24 1.1.1.1 0 100 0 600 i
>i 172.168.25.0/24 1.1.1.1 0 100 0 ?
  Announced routes count = 0
```

TWAMP over L3VPN Configuration

This chapter contains a complete sample TWAMP over L3VPN configuration. Two-way Active Measurement Protocol (TWAMP) is an open protocol for measuring network performance between any two devices. The TWAMP MPLS transport is implemented as part of supporting TWAMP on routers which acts as MPLS routers both in the roles of LERs as well as intermediate routers. OcNOS version 6.0.0 also supports the end to end statistics calculation when multiple paths are available between sender and reflector with multihop support.

The user can use the link delay metrics such as average, minimum, and maximum delay, and delay variance to determine the network latency. Using link delay metrics will enable troubleshooting latency issues or apply Traffic Engineering (TE) solutions to meet Service Level Agreements (SLAs).

The TWAMP protocol is designed to do such measurements, and a basic implementation of this protocol has already been implemented in OcNOS. This feature here is a TWAMP protocol in OcNOS where the focus will be on accuracy and configurable advertisement of the measured data.

L3VPN (based on MPLS) Supported scenarios:

In general, TWAMP over L3VPN works on:

- CE-CE Overlay Only
- CE-PE Overlay Only
- PE-PE Both Under lay and over lay.

Topology

Displays a sample TWAMP over L3VPN topology.

- CE1 and CE2 are customer edge routers
- PE1 and PE2 are IPv4 Provider Edge routers
- P1 is the router at the core of the IPv4 MPLS provider network.




Figure 99. TWAMP over L3VPN Topology

Configure TWAMP over L3VPN for LDP

PE1

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)# ip vrf vrf100 | Create a new VRF named vrf100 |
| (config-vrf)#rd 100:1 | Assign the route distinguisher (RD) value as 100:1 |
| (config-vrf)#route-target both 100:1 | Import routes between route target (RT) ext-communities 100 and 1 |
| (config-vrf)#exit | Exit VRF mode |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 18.18.18.18/32 secondary | Assign IP address to Loopback interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe8 | Enter Interface mode |
| (config-if)# ip address 10.1.1.18/24 | Assign IP address to interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xel1 | Enter Interface mode |
| (config-if)# ip vrf forwarding vrf100 | Bind the interface connected to the CE1 router with VRF 100 |
| (config-if)# ip address 100.1.1.1/24 | Assign IP address to interface |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 18.18.18.18 | Configure router id as loopback address |
| (config-router)#transport-address ipv4 18.18.18.18 | Configure ldp transport address as loopback address |
| (config-router)#exit | Exit from the router ldp mode |
| (config)# router ospf 1 | Enter Router OSPF mode |
| (config-router)# ospf router-id 18.18.18.18 | Configure OSPF router-id |
| (config-router)# network 10.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# network 18.18.18.18/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# commit | Commit the configurations |
| (config-router)# exit | Exit from router OSPF mode |
| (config)# router bgp 100 | Enter BGP router mode |
| (config-router)# bgp router-id 18.18.18.18 | Configure BGP router-id |

| | |
|---|---|
| (config-router)# neighbor 8.8.8.8 remote-as 100 | Configure PE2 as an iBGP4+ neighbor |
| (config-router)# neighbor 8.8.8.8 up-date-source lo | Update the source as loopback for iBGP peering with the remote PE2 router |
| (config-router)# address-family vpnv4 unicast | Enter address-family vpnv4 mode |
| (config-router-af)# neighbor 8.8.8.8 activate | Activate the PE2 neighbor in the vpnv4 address family |
| (config-router-af)# neighbor 8.8.8.8 next-hop-self | Activate the neighbor as next hop self |
| (config-router-af)#exit | Exit form address family |
| (config-router)# address-family ipv4 vrf vrf100 | Enter the IPv4 address family for VRF 100 |
| (config-router-af)# redistribute connected | Redistribute connected routes |
| (config-router-af)# neighbor 100.1.1.2 re-mote-as 200 | Configure CE1 neighbor in the vrf address family |
| (config-router-af)# neighbor 100.1.1.2 acti-vate | Activate the CE1 neighbor |
| (config-router-af)#exit | Exit form address family |
| (config-router)# commit | Commit the configurations |

P1

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 3.3.3.3/32 secondary | Assign IP address to Loopback interface |
| (config-if)#exit | Exit Interface mode |
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 3.3.3.3 | Configure router id as loopback address |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure ldp transport address as loopback address |
| (config-router)#exit | Exit from the router ldp mode |
| (config)#interface xe14 | Enter Interface mode |
| (config-if)# ip address 10.1.1.3/24 | Assign IP address to interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe15 | Enter Interface mode |
| (config-if)# ip address 11.1.1.3/24 | Assign IP address to interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit interface mode |
| (config)# router ospf 1 | Enter Router OSPF mode |
| (config-router)# ospf router-id 18.18.18.18 | Configure OSPF router-id |

| | |
|---|--|
| (config-router)# network 10.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# network 3.3.3.3/32 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# network 11.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# commit | Commit the configurations |
| (config-router)# exit | Exit from router OSPF mode |

PE2

| | |
|---|---|
| #configure terminal | Enter Configure mode. |
| (config)#ip vrf vrf100 | Create a new VRF named vrf100 |
| (config-vrf)#rd 101:1 | Assign the route distinguisher (RD) value as 101:1 |
| (config-vrf)#route-target both 100:1 | Import routes between route target (RT) ext-communities 100 and 1 |
| (config-vrf)#exit | Exit VRF mode |
| (config)#router ldp | Enter the Router LDP mode |
| (config-router)#router-id 8.8.8.8 | Configure router id as loopback address |
| (config-router)#transport-address ipv4 8.8.8.8 | Configure ldp transport address as loopback address |
| (config-router)#exit | Exit from the router ldp mode |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 8.8.8.8/32 secondary | Assign IP address to Loopback interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe24 | Enter Interface mode |
| (config-if)# ip address 11.1.1.8/24 | Assign IP address to interface |
| (config-if)#enable-ldp ipv4 | Enable LDP on the physical interface |
| (config-if)#label-switching | Enable label switching on the interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe25 | Enter Interface mode |
| (config-if)# ip vrf forwarding vrf100 | Bind the interface connected to the CE1 router with VRF 100 |
| (config-if)# ip address 101.1.1.1/24 | Assign IP address to interface |
| (config-if)#exit | Exit interface mode |
| (config)# router ospf 1 | Enter Router OSPF mode |
| (config-router)# ospf router-id 8.8.8.8 | Configure OSPF router-id |
| (config-router)# network 11.1.1.0/24 area 0.0.0.0 | Define the network on which OSPF runs and associate area id. |
| (config-router)# network 8.8.8.8/32 area 0.0.0.0 | Define the network on which OSPF runs and |

| | |
|---|---|
| | associate area id. |
| (config-router)#exit | Exit from router OSPF mode |
| (config)# router bgp 100 | Enter BGP router mode |
| (config-router)# bgp router-id 8.8.8.8 | Configure BGP router-id |
| (config-router)# neighbor 18.18.18.18 re-mote-as 100 | Configure PE1 as an iBGP4+ neighbor |
| (config-router)# neighbor 18.18.18.18 up-date-source lo | Update the source as loopback for iBGP peering with the remote PE1 router |
| (config-router)# address-family vpnv4 unicast | Enter address-family vpnv4 mode |
| (config-router-af)# neighbor 18.18.18.18 ac-tivate | Activate the PE1 neighbor in the vpnv4 address family |
| (config-router-af)# neighbor 18.18.18.18 next-hop-self | Activate the neighbor as next hop self |
| (config-router-af)#exit | Exit form address family |
| (config-router)# address-family ipv4 vrf vrf100 | Enter the IPv4 address family for VRF 100 |
| (config-router-af)# redistribute connected | Redistribute connected routes |
| (config-router-af)# neighbor 101.1.1.2 re-mote-as 200 | Configure CE2 neighbor in the vrf address family |
| (config-router-af)# neighbor 101.1.1.2 acti-vate | Activate the CE2 neighbor |
| (config-router-af)#exit | Exit form address family |
| (config-router)# commit | Commit the configurations |

CE1

| | |
|--|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 37.37.37.37/32 secondary | Assign IP address to Loopback interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe24 | Enter Interface mode |
| (config-if)# ip address 100.1.1.2/24 | Assign IP address to interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe26 | Enter Interface mode |
| (config-if)# ip address 200.1.1.1/24 | Assign IP address to interface |
| (config-if)#exit | Exit interface mode |
| (config)# router bgp 200 | Enter BGP router mode |
| (config-router)# bgp router-id 37.37.37.37 | Configure BGP router-id |
| (config-router)#neighbor 100.1.1.1 remote-as 100 | Configure PE1 as an eBGP4+ neighbor |
| (config-router)# address-family ipv4 unicast | Enter address-family IPv4 unicast mode |
| (config-router-af)# redistribute connected | Redistribute the connected route under address family IPv4 unicast |

| | |
|---|--|
| (config-router-af)# neighbor 100.1.1.1 activate | Activate the neighbor in the IPv4 address family |
| (config-router-af)#exit | Exit form address family |
| (config-router)# commit | Commit the configurations |

CE2

| | |
|--|--|
| #configure terminal | Enter Configure mode. |
| (config)#interface lo | Enter loopback interface mode |
| (config-if)#ip address 2.2.2.2/32 secondary | Assign IP address to Loopback interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe14 | Enter Interface mode |
| (config-if)# ip address 101.1.1.2/24 | Assign IP address to interface |
| (config-if)#exit | Exit Interface mode |
| (config)#interface xe15 | Enter Interface mode |
| (config-if)# ip address 201.1.1.1/24 | Assign IP address to interface |
| (config-if)#exit | Exit interface mode |
| (config)# router bgp 300 | Enter BGP router mode |
| (config-router)# bgp router-id 2.2.2.2 | Configure BGP router-id |
| (config-router)#neighbor 101.1.1.1 remote-as 100 | Configure PE2 as an eBGP4+ neighbor |
| (config-router)# address-family ipv4 unicast | Enter address-family IPv4 unicast mode |
| (config-router-af)# redistribute connected | Redistribute the connected route under address family IPv4 unicast |
| (config-router-af)# neighbor 101.1.1.1 acti-vate | Activate the neighbor in the IPv4 address family |
| (config-router-af)#exit | Exit form address family |
| (config-router)# commit | Commit the configurations |

TWAMP Configuration on Sender (PE1)

TWAMP sender is configured to measure the delay on interface Loopback on PE1

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)# hardware-profile filter twamp-ipv4 enable | Enable hardware filter for ipv4 to configure TWAMP meas-urement configs |
| (config)#commit | Commit the configuration |
| (config)# twamp-light control | Enable TWAMP light controller on PE1 |
| (config-twamp-light-con)# control-admin-state enable | Enable TWAMP Controller admin state |
| (config)#interface lo | Enter Interface Loopback mode |
| (config-if)# delay-measurement dynamic twamp reflector-ip 11.1.1.8 | Configure delay measurement on interface Loopback to re-lector PE2 |

| | |
|--------------------|---------------------------|
| (config-if)#commit | Commit the configurations |
| (config-if)#end | Return to privilege mode |

TWAMP Configuration on Reflector (PE2)

Configure TWAMP Reflector as interface xe24 on PE2 (Towards core)

| | |
|--|---|
| #configure terminal | Enter Configure mode. |
| (config)# hardware-profile filter twamp-ipv4 enable | Enable hardware filter for ipv4 to configure TWAMP meas-urement configs |
| (config)#commit | Commit the configuration |
| (config)# twamp-light reflector | Enable TWAMP light Reflector on PE2 |
| (config-twamp-light-ref)# reflec-tor-admin-state enable | Enable the TWAMP reflector admin state |
| (config-twamp-light-ref)# reflector-name pe2 reflector-ip ipv4 11.1.1.8 | Configure TWAMP reflector IP as PE2 interface IP |
| (config-twamp-light-ref)#commit | Commit the configurations |
| (config-if)#end | Return to privilege mode |

Validation

1. Verify ping from PE1 to PE2

```
PE1#ping 11.1.1.8
Press CTRL+C to exit
PING 11.1.1.8 (11.1.1.8) 56(84) bytes of data.
64 bytes from 11.1.1.8: icmp_seq=1 ttl=63 time=0.683 ms
64 bytes from 11.1.1.8: icmp_seq=2 ttl=63 time=0.491 ms
64 bytes from 11.1.1.8: icmp_seq=3 ttl=63 time=0.594 ms
```

2. Verify mpls l3vpn ping from PE1

```
PE1#ping mpls l3vpn vrf100 101.1.1.0/24
Sending 5 MPLS Echos to 101.1.1.0, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

!
!
!
!
!

Success Rate is 100.00 percent (5/5)
```

3. Verify the TWAMP statistics on all the configured interfaces on PE1

In the below verification command, packets sent and received showing as equal. So all the TWAMP packets received reply for all the sent packets for the delay measurement. Showing all the Round Trip Delay and Reverse Delay timers.

```
PE1#sh twamp-statistics
=====
TWAMP Test-Session Statistics
=====
Test-Session Name      : __internal_interface_xe8
Start Time             : 2019 Feb 14 15:21:58
Elapsed time(milli sec) : 6003
Packets Sent           : 6
Packets Received       : 6
Packet Loss(%)         : 0.00
Round Trip Delay(usec)
  Minimum               : 17
  Maximum               : 19
  Average               : 18
Forward Delay(usec)
  Minimum               : (*)
  Maximum               : (*)
  Average               : (*)
Reverse Delay(usec)
  Minimum               : (*)
  Maximum               : (*)
  Average               : (*)
Round Trip Delay Variation(usec)
  Minimum               : 16
  Maximum               : 19
  Average               : 17
Forward Delay Variation(usec)
  Minimum               : (*)
  Maximum               : (*)
  Average               : (*)
Reverse Delay Variation(usec)
  Minimum               : (*)
  Maximum               : (*)
  Average               : (*)

(*) - Time is not in sync between Sender and Reflector
```

4. Verify the List of all interfaces that are currently participating in Delay measurement

```
PE1#sh twamp-statistics interfaces
Interface Last Advertisement Delay(us) Min(us) Max(us) Var(us) Loss(%)
Xe8      2019-02-14 15:22:00 10      10      10      0      Not Enabled
```

5. Verify the Detailed list of TWAMP delay measurement information on interface xe8

```
PE1#sh twamp-statistics interfaces xe8
Interface name      : xe8
Sender IP          : 10.1.1.18
Reflector IP       : 11.1.1.8
Reflector port     : 862
DSCP value         : 0
HW Status          : HW rules installed
Last Advertised stats:
Time: 2019-02-14 15:22:00
Average delay      : 10
Minimum delay      : 10
Maximum delay      : 10
Average delay variation: 0
Minimum delay variation: 0
Maximum delay variation: 0
Packets sent       : 1
Packets received   : 1
Packets timeout    : 0
```

```
Packet Loss: Not Enabled
Last Calculated stats:
Time: 2019-02-14 15:22:00
Average delay      : 10
Minimum delay      : 10
Maximum delay      : 10
Average delay variation: 0
Minimum delay variation: 0
Maximum delay variation: 0
Packets sent       : 1
Packets received    : 1
Packets timeout     : 0
Packet Loss : Not Enabled
```

Bridge Virtual Interface (BVI) Over L3VPN

Overview

A Bridge Virtual Interface (BVI) is a virtual interface on a router that acts like a routed interface and is associated with a single bridge domain.

BVI Interface acts as L3 routed interface gateway between bridge domain and L3VPN for traffic exchange. The incoming tagged packet from the L2 sub-interface consolidated itself into a bridge domain. The bridge domain in turn uses the BVI interface to forward the IP traffic to the L3VPN tunnel.

For details on BVI Over 6vPE, see [Bridge Virtual Interface \(BVI\) Over 6vPE \(page 2621\)](#). For QoS and ACL configuration for BVIs, see the *QoS and ACL over Bridge Virtual Interface (BVI)* chapter in Layer 3 guide.

Characteristics of BVI Over L3VPN

- The BVI functions as an L3 routed interface for a bridge domain, allowing IP traffic from L2 subinterfaces within the domain to be routed to L3VPN tunnels.
- L2 subinterfaces can be grouped into a bridge domain under the BVI. The BVI aggregates traffic from multiple subinterfaces within the same domain and routes it as required.
- The BVI remains operationally “up” as long as at least one of the subinterfaces in the bridge domain is active.
- BVI supports unicast forwarding of IPv4 traffic between other L3 interfaces and L3VPNs. The BVI can function as a DHCP server or relay for IP address allocation within the network.

Benefits

- The BVI enables seamless communication between L2 bridge domains and L3 networks (e.g., L3VPN), allowing for flexible traffic forwarding between the two layers.
- The BVI serves as an L3 gateway for M-plane traffic, routing it efficiently to L3VPN tunnels for further processing or external routing.
- The BVI remains operational even if only one subinterface is active, ensuring high availability and fault tolerance.

Limitations

- Only Q1 and Q2-based platforms are supported.
- Everything related to ipv6 is not supported.
- L2 sub interfaces support only pop and pop2-tag vlan rewrite translations and for the purpose of bridge-domain support, all member L2 sub interfaces should have relevant rewrite configurations to make sure that all traffic for bridge-domain has uniform encapsulation.
- The dot1q and dot1ad encapsulations with range are not supported for the BVI.
- BVI is a generic L3 interface and allows users to configure any / all existing OcNOS features. However, only a few of these are supported in release 6.6.0 and those are IP address related, VRF, any routing protocol specific commands (OSPF/IS-IS/BGP related), MTU, DHCP server/relay, Ingress/Egress ACL and QoS commands related to marking.

- L3 ACL applied at BVI is only relevant for routed traffic. Bridged traffic between L2 sub-interfaces will not be subjected to L3 ACL configured at BVI.
- No Interface counters are supported for BVI Interface.
- BVI can't be used as a network interface for all transports in MPLS core network.
- For QoS, only marking related CLIs are supported. Policing, rate limiting, shaping and other queuing features are not supported at BVI level. However, existing QoS scheduling and queuing features on other interfaces can be used in conjunction with the BVI interface.
- QoS features like shaping, queueing, policing, re-marking are not supported on the BVI interface associated to VPLS.
- BUM traffic is not supported.

Prerequisites

- Define Interfaces and Loopback Addresses:

Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish efficient network routing and communication.

```
interface lo
 ip address 127.0.0.1/8
 ip address 135.1.1.27/32 secondary
 ipv6 address ::1/128

interface po6
 ip address 10.1.1.1/30
interface xe6
 channel-group 6 mode active
```

- Configure IGP for Dynamic Routing:

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

- ISIS Configuration:

```
router isis 1
 is-type level-2-only
 metric-style wide
 mpls traffic-eng router-id 10.12.183.1
 mpls traffic-eng level-2
 capability cspf
 dynamic-hostname
 fast-reroute ti-lfa level-2 proto ipv4
 net 49.0000.0000.0027.00
 passive-interface lo

interface po6
 isis network point-to-point
 ip router isis 1
```

- OSPF Configuration:

```
router ospf 100
 ospf router-id 10.12.183.1
 network 10.12.183.1/32 area 0.0.0.0
 network 10.1.1.0/24 area 0.0.0.0
 network 10.1.1.0/24 area 0.0.0.0
```

- Configure LDP for Label Transport:
Configure label-switching on all the nodes to help exchange the network packets at a lower lever rather than the traditional network layer and enable ldp to transport those labels.

```
router ldp
targeted-peer ipv4 10.12.183.3
exit-targeted-peer
transport address ipb4 10.12.183.1

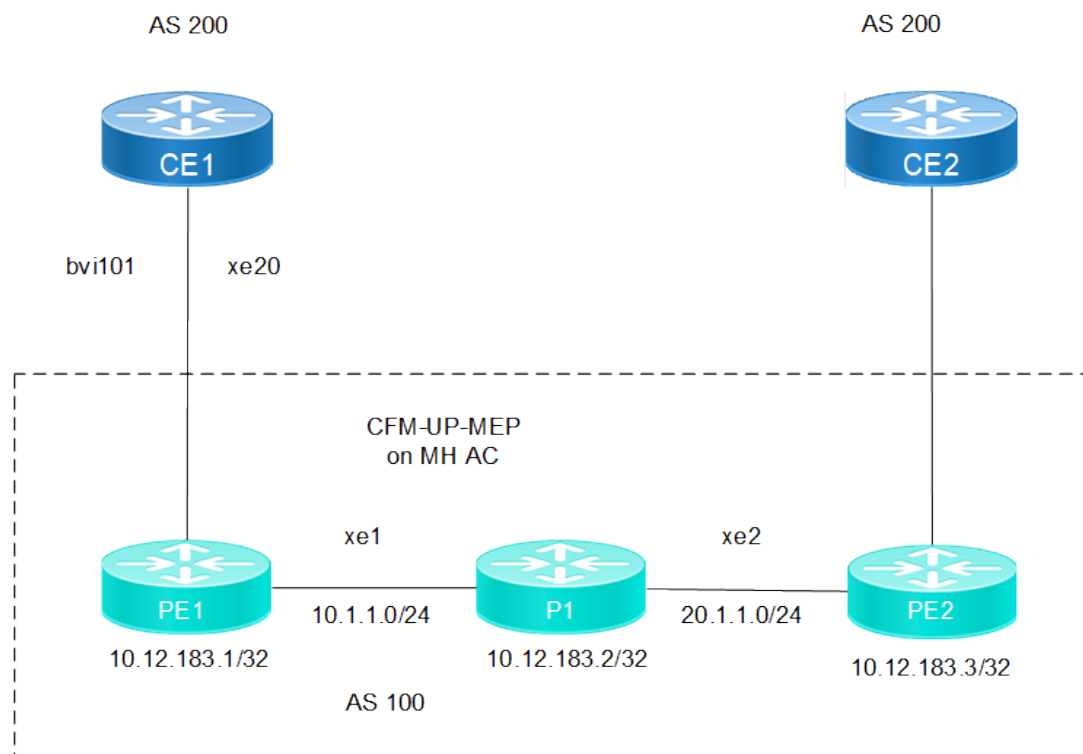
Interface po6
label-switching
enable-ldp ipv4
```

Configuration

Topology

The sample topology includes Edge nodes (PE1 and PE2), core Node (P1), and Customer Edge nodes (CE1 and CE2).

Figure 100. BVI Interface Topology



Note: Before configuring BVI-L3VPN, meet all the pre-requisite for the following nodes:

- Edge nodes: PE1 and PE2
- Core nodes: P1, P3, and P4

Configure CE1 Router

1. Set up the interface with VLAN encapsulation and assign an IP address.

```
CE1(config)#interface xe20.101
CE1(config-if)#encapsulation dot1q 101
CE1(config-if)#ip address 101.0.1.2/24
CE1(config-if)#exit
```

2. with a neighbor in the same AS (200).

```
CE1(config)# router bgp 200
CE1(config-router)#neighbor 101.0.1.1 remote-as 100
CE1(config-router)#address-family ipv4 unicast
CE1(config-router-af)#neighbor 101.0.1.1 activate
CE1(config-router-af)#exit-address-family
CE1(config-router)#exit
```

Configure PE1 Router

1. Create a VRF instance (VRF101) and define route targets.

```
PE1(config)# ip vrf VRF101
PE1(config-vrf)# rd 10.12.183.1:100
PE1(config-vrf)# route-target both 100:101
PE1(config-vrf)# exit
```

2. Configure the access port with VLAN encapsulation and enable rewriting.

```
PE1(config)#interface xe20.101 switchport
PE1(config-if)#encapsulation dot1q 101
PE1(config-if)#rewrite pop
PE1(config-if)# exit
```

3. Set up a BVI for VRF forwarding.

```
PE1(config)# interface bvi101
PE1(config-if)#ip vrf forwarding VRF101
PE1(config-if)#ip address 101.0.1.1/24
PE1(config-if)# exit
```

4. Configure BGP with internal (PE) and external (CE) neighbors.

PE1(config)# router bgp 100

```
PE1(config)# router bgp 100
PE1(config-router)# neighbor 10.12.183.3 remote-as 100
PE1(config-router)# neighbor 10.12.183.3 update-source lo
PE1(config-router)#address-family ipv4 unicast
PE1(config-router-af)#neighbor 10.12.183.3 activate
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family vpnv4 unicast
PE1(config-router-af)#neighbor 10.12.183.3 activate
PE1(config-router-af)#exit-address-family
```

5. Enable VPNv4 and redistribute connected routes.

```
PE1(config-router)#address-family ipv4 vrf VRF101
PE1(config-router-af)#redistribute connected
PE1(config-router-af)#neighbor 101.0.1.2 remote-as 200
PE1(config-router-af)#neighbor 101.0.1.2 activate
PE1(config-router-af)#exit-address-family
PE1(config-router)#exit
```

6. Set up a bridge domain to associate the interface with BVI.

```
PE1(config)# bridge-domain 101
PE1(config)#interface xe20.101
PE1(config)#routed-interface bvi101
```

Configure PE2 Router

1. Create the same VRF instance (VRF101) with route targets.

```
PE2(config)# ip vrf VRF101
PE2(config-vrf)# rd 10.12.183.3:100
PE2(config-vrf)# route-target both 100:101
PE2(config-vrf)#exit
```

2. Configure the access port with VLAN encapsulation and VRF forwarding.

```
PE2(config)# interface xe30.101
PE2(config-if)#encapsulation dot1q 101
PE2(config-if)#ip vrf forwarding VRF101
PE2(config-if)#ip address 103.0.1.1/24
PE2(config-f)# exit
```

3. Assign an IP address to the interface.

```
PE2(config)# router bgp 100
PE2(config-router)# neighbor 10.12.183.1 remote-as 100
PE2(config-router)# neighbor 10.12.183.1 update-source lo
PE2(config-router)#address-family ipv4 unicast
PE2(config-router-af)#neighbor 10.12.183.1 activate
PE2(config-router-af)#exit-address-family
```

4. Configure BGP with internal (PE) and external (CE) neighbors.

```
PE2(config-router)#address-family vpnv4 unicast
PE2(config-router-af)#neighbor 10.12.183.1 activate
PE2(config-router-af)#exit-address-family
```

5. Enable VPNv4 and redistribute connected routes.

```
PE2(config-router)#address-family ipv4 vrf VRF101
PE2(config-router-af)#redistribute connected
PE2(config-router-af)#neighbor 103.0.1.2 remote-as 200
PE2(config-router-af)#neighbor 103.0.1.2 activate
PE2(config-router-af)#exit-address-family
PE2(config-router)#exit
```

Configure CE2 Router

1. Set up the interface with VLAN encapsulation and assign an IP address.

```
CE2#configure terminal
CE2(config)#interface xe30.101
CE2(config-if)#encapsulation dot1q 101
CE2(config-if)#ip address 103.0.1.2/24
CE2(config-if)#exit
```

2. Configure BGP with a neighbor in the same AS (200).

```
CE2(config)# router bgp 200
CE2(config-router)#neighbor 103.0.1.1 remote-as 100
CE2(config-router)#address-family ipv4 unicast
CE2(config-router-af)#neighbor 103.0.1.1 activate
CE2(config-router-af)#exit-address-family
CE2(config-router)#exit
```

Running Configuration

CE1

```
#show running-config
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
!
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics ingress-acl enable
!
qos enable
!
hostname CE1
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
!
ip vrf management
!
interface ce0
!
interface ce1
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128

interface xe20
!
interface xe20.101
 encapsulation dot1q 101
 ip address 101.0.1.2/24
!
```

```
router bgp 200
 neighbor 101.0.1.1 remote-as 100
 !
 address-family ipv4 unicast
 neighbor 101.0.1.1 activate
 exit-address-family
 !
 exit
 !
 !
end
```

PE1

```
#show running-config
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
!
service password-encryption
!
logging console 5
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile filter ingress-ipv4-ext enable
hardware-profile filter egress-ipv4-ext enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
qos enable
!
hostname PE1
port ce2 breakout 4X10g
no ip domain-lookup
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
bridge 1 protocol ieee vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server community test vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
feature rsyslog
logging remote server 10.12.100.252 5 port 1514 vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
router-id 10.12.183.1
```

```
!  
ip vrf management  
!  
ip vrf VRF101  
  rd 10.12.183.1:100  
  route-target both 100:101  
!  
router ldp  
  fast-reroute  
  session-protection duration 40  
  targeted-peer ipv4 10.12.183.3  
  exit-targeted-peer-mode  
  transport-address ipv4 10.12.183.1  
!  
router rsvp  
!  
interface po6  
  ip address 10.1.1.1/24  
  label-switching  
  ip router isis ISIS-IGP-100  
  enable-ldp ipv4  
!  
interface bv101  
  ip vrf forwarding VRF101  
  ip address 101.0.1.1/24  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 10.12.183.1/32 secondary  
  ipv6 address ::1/128  
  ip router isis ISIS-IGP-100  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xel  
  channel-group 6 mode active  
!  
interface xe20  
!  
interface xe20.101 switchport  
  encapsulation dot1q 101  
  rewrite pop  
!  
  
  exit  
!  
router ospf 100  
  ospf router-id 10.12.183.1  
  network 10.1.1.0/24 area 0.0.0.0  
  network 10.12.183.1/32 area 0.0.0.0  
!  
router isis ISIS-IGP-100  
  is-type level-1  
  metric-style wide  
  mpls traffic-eng router-id 10.12.183.1  
  mpls traffic-eng level-1  
  capability cspf  
  dynamic-hostname  
  fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp  
  bfd all-interfaces  
  net 49.0001.0000.0000.0001.00
```

```

    passive-interface lo
    !
router bgp 100
  neighbor 10.12.183.3 remote-as 100
  neighbor 10.12.183.3 update-source lo
  !
  address-family ipv4 unicast
  neighbor 10.12.183.3 activate
  exit-address-family
  !
  address-family vpnv4 unicast
  neighbor 10.12.183.3 activate
  exit-address-family
  !
  address-family ipv4 vrf VRF101
  redistribute connected
  neighbor 101.0.1.2 remote-as 200
  neighbor 101.0.1.2 activate
  exit-address-family
  !
  exit
  !
bridge-domain 100
  interface xe20.101
  routed-interface bvi101
  !
  !
end

```

P1

```

#show running-config
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
!
service password-encryption
!
logging console 5
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
qos enable
!
hostname P1
no ip domain-lookup
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
bridge 1 protocol ieee vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server community test vrf management
feature dns relay
ip dns relay
ipv6 dns relay

```



```
feature ntp vrf management
ntp enable vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
!
router-id 10.12.183.2
!
ip vrf management
!
router ldp
  transport-address ipv4 10.12.183.2
!
router rsvp
!
interface po6
  ip address 10.1.1.2/24
  label-switching
  ip router isis ISIS-IGP-100
  enable-ldp ipv4
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface lo
  ip address 127.0.0.1/8
  ip address 10.12.183.2/32 secondary
  ipv6 address ::1/128
  ip router isis ISIS-IGP-100
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
interface xe1
  channel-group 6 mode active
!
interface xe2
  ip address 20.1.1.1/24
  label-switching
  ip router isis ISIS-IGP-100
  enable-ldp ipv4
!
exit
!
router ospf 100
  ospf router-id 10.12.183.2
  network 10.1.1.0/24 area 0.0.0.0
  network 10.12.183.2/32 area 0.0.0.0
  network 20.1.1.0/24 area 0.0.0.0
!
router isis ISIS-IGP-100
  is-type level-1
  metric-style wide
  mpls traffic-eng router-id 10.12.183.2
  mpls traffic-eng level-1
  capability cspf
  dynamic-hostname
  fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
  bfd all-interfaces
  net 49.0001.0000.0000.0002.00
  passive-interface lo
```

```
!  
end
```

PE2

```
#show running-config  
!  
feature netconf-ssh vrf management  
feature netconf-tls vrf management  
no feature netconf-ssh  
no feature netconf-tls  
!  
service password-encryption  
!  
logging console 5  
snmp-server enable traps link linkDown  
snmp-server enable traps link linkUp  
!  
hardware-profile statistics ingress-acl enable  
hardware-profile statistics ac-lif enable  
!  
qos enable  
!  
hostname PE2  
no ip domain-lookup  
ip domain-lookup vrf management  
ip name-server vrf management 10.12.3.23  
bridge 1 protocol ieee vlan-bridge  
tfo Disable  
errdisable cause stp-bpdu-guard  
no feature telnet vrf management  
no feature telnet  
feature ssh vrf management  
no feature ssh  
snmp-server enable snmp vrf management  
snmp-server view all .1 included vrf management  
snmp-server community test vrf management  
feature dns relay  
ip dns relay  
ipv6 dns relay  
feature ntp vrf management  
ntp enable vrf management  
lldp run  
lldp tlv-select basic-mgmt port-description  
lldp tlv-select basic-mgmt system-name  
lldp tlv-select basic-mgmt system-capabilities  
lldp tlv-select basic-mgmt system-description  
lldp tlv-select basic-mgmt management-address  
lldp notification-interval 10  
!  
router-id 10.12.183.3  
!  
ip vrf management  
!  
ip vrf VRF101  
  rd 10.12.183.3:100  
  route-target both 100:101  
!  
router ldp  
  fast-reroute  
  session-protection duration 40  
  targeted-peer ipv4 10.12.183.1  
  exit-targeted-peer-mode  
  transport-address ipv4 10.12.183.3  
!  
interface eth0
```

```
ip vrf forwarding management
ip address dhcp
!
interface lo
ip address 127.0.0.1/8
ip address 10.12.183.3/32 secondary
ipv6 address ::1/128
ip router isis ISIS-IGP-100
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface xe2
ip address 20.1.1.2/24
label-switching
ip router isis ISIS-IGP-100
enable-ldp ipv4
!
interface xe30
!
interface xe30.101
description L3VPN-VRF101
encapsulation dot1q 101
ip vrf forwarding VRF101
ip address 103.0.1.1/24
!
exit
!
router ospf 100
ospf router-id 10.12.183.3
network 10.12.183.3/32 area 0.0.0.0
network 20.1.1.0/24 area 0.0.0.0
!
router isis ISIS-IGP-100
is-type level-1
metric-style wide
mpls traffic-eng router-id 10.12.183.3
mpls traffic-eng level-1
capability cspf
dynamic-hostname
fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
bfd all-interfaces
net 49.0001.0000.0000.0003.00
passive-interface lo
!
router bgp 100
neighbor 10.12.183.1 remote-as 100
neighbor 10.12.183.1 update-source lo
!
address-family ipv4 unicast
neighbor 10.12.183.1 activate
exit-address-family
!
address-family vpnv4 unicast
neighbor 10.12.183.1 activate
exit-address-family
!
address-family ipv4 vrf VRF101
redistribute connected
neighbor 103.0.1.2 remote-as 200
neighbor 103.0.1.2 activate
exit-address-family
!
exit
!
end
```

CE2

```
#show running-config
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
!
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics ingress-acl enable
!
qos enable
!
hostname CE2
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
!
ip vrf management
!

interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128

interface xe30
!
interface xe30.101
 encapsulation dot1q 101
 ip address 103.0.1.2/24
!
router bgp 200
 neighbor 103.0.1.1 remote-as 100
!
 address-family ipv4 unicast
 neighbor 103.0.1.1 activate
```

```

exit-address-family
!
exit
!
!
end

```

Validation

PE1

To Verify the L2 interfaces and BVI interface attached on a bridge domain:

```

PE1#show running-config bridge-domain
bridge-domain 101
Interface xe20.101
routed-interface bvi101

```

To Verify the link status of L2 Interfaces attached to bridge-domain:

```

PE1#show bridge-domain

Total number of bridge-domains Configured: 1

Bridge Id  interfaces      Status
-----+-----+-----
101        xe20.101         UP

```

To Verify the link status of the BVI interface:

```

PE1#show ip interface bvi101 brief

'*' - address is assigned by dhcp client

Interface      IP-Address      Admin-Status      Link-Status
bvi101         101.0.1.1       up                 up

```

To Verify BGP session between PE1-CE1

```

PE1#show ip bgp summary
BGP router identifier 10.12.183.1, local AS number 100
BGP table version is 11
1 BGP AS-PATH entries
0 BGP community entries

Neighbor      V AS  MsgRcv MsgSen TblVer InQ OutQ Up/Down  State/PfxRcd  Desc
10.12.183.3  4 100  280    283    11    0   0  00:18:47    0

Total number of neighbors 1

Total number of Established sessions 1
BGP router identifier 101.0.1.1, local AS number 100
BGP VRF VRF101 Route Distinguisher: 10.12.183.1:100
BGP table version is 1
1 BGP AS-PATH entries
0 BGP community entries

Neighbor      V AS  MsgRcv MsgSen TblVer InQ OutQ Up/Down  State/PfxRcd  Desc
101.0.1.2     4 200   10     13     1     0   0  00:03:54    0

Total number of neighbors 1

Total number of Established sessions 1

```

To Verify the route between PE1-CE1

```
PE1#show ip route vrf VRF101
IP Route Table for VRF "VRF101"
C          101.0.1.0/24 is directly connected, bvi101, installed 01:59:42, last update 01:59:42 ago
C          127.0.0.0/8 is directly connected, lo.VRF101, installed 01:59:46, last update 01:59:46 ago
```

To Verify VRF Ping between PE1-CE1

```
PE1#ping ip vrf VRF 101.0.1.2 vrf VRF101
Press CTRL+C to exit
PING 101.0.200.2 (101.0.200.2) 100(128) bytes of data.
108 bytes from 101.0.200.2: icmp_seq=1 ttl=64 time=0.432 ms
108 bytes from 101.0.200.2: icmp_seq=2 ttl=64 time=0.427 ms
108 bytes from 101.0.200.2: icmp_seq=3 ttl=64 time=0.348 ms

--- 101.0.200.2 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2070ms
rtt min/avg/max/mdev = 0.348/0.402/0.432/0.038 ms
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms/Acronym | Description |
|-------------------|--|
| BUM | Broadcast, Unknown, Multicast |
| BVI | Bridge Virtual Interface |
| L3VPN | Layer 3 Virtual Private Network |
| IGP | Interior Gateway Protocol |
| ISIS | Intermediate System to Intermediate System |
| OSPF | Open Shortest Path First |
| LDP | Label Distribution Protocol |

DSCP Preserve

Layer 3 VPN services use the DSCP field in the IP header to convey QoS information, determining packet priority and treatment (e.g., DSCP 46 for high-priority voice traffic). When L3 VPN traffic enters the MPLS domain, QoS is managed by the EXP field in the MPLS label, typically derived from the DSCP value. In Qumran-based systems, the EXP field is updated based on the remarked DSCP value. By default, remarking the DSCP at the access interface affects both the IP header DSCP field and the MPLS EXP field. Service managers need control over whether to remark the DSCP or not while configuring EXP field updates.

Network operator can control preserving DSCP in the edge routers at global level or per VRF level or combination of both. To provide more control, the following DSCP preservation settings are available:

- [DSCP Preserve Global \(page 2556\)](#) - When DSCP preserve is enabled globally, it applies to all VRFs, preserving the DSCP value even if remarking is configured.
- [DSCP Preserve Per-VRF \(page 2570\)](#) - Per-VRF configuration allows DSCP preservation only for specific VPN services, offering more granular control for network administrators.

DSCP Preserve Behavior: Below table shows the way DSCP preserve global and per VRF configuration impact the DSCP value of the traffic received on the VRF of interest when dscp-to-queue profile is updated with remark DSCP option.

Table 88.

| DSCP-preserve | per-vrf not configured | per-vrf enabled | per-vrf disabled |
|----------------|------------------------|-----------------|------------------|
| Global enable | preserved | preserved | not preserved |
| Global disable | not preserved | preserved | not preserved |

Network administrators can further refine control using global or custom DSCP-to-queue profiles and policy-map configurations to remark DSCP values based on specific match criteria. This provides flexible and tailored QoS management across different VPN services.

DSCP Preserve Characteristics

- QoS information mapping ensures that the QoS information carried in the DSCP field of the IP header is correctly mapped to the MPLS EXP field.
- DSCP remarking controls whether the DSCP value in the IP header is remarked at the access interface, affecting both customer traffic and the MPLS domain.
- In hardware like Qumran, the EXP field update is based on the remarked DSCP value, requiring configuration to manage how and if DSCP values are modified.
- Ingress configuration requirement introduces the need for ingress-side configuration to properly align the DSCP and EXP field values, ensuring correct QoS handling.
- DSCP preserve provides service managers with control over whether DSCP values are modified or preserved during the process of mapping to MPLS headers.

DSCP Preserve Global

This method of configuration allows Network operator to enable or disable DSCP preserve for all the services in an edge router.

Prerequisites

The L3VPN service will operate seamlessly from end to end without any packet loss, using either OSPF or IS-IS as the IGP, with traffic profiles configured based on user-defined DSCP values.

- Globally enable dscp-preserve using the `mpls lsp-encap-dscp-preserve` command.
- Global level configuration for PE1:

```
config t
qos enable
!
qos profile dscp-to-queue profile1
dscp 24 queue 3 dscp 32
```

- interface level configuration for PE1:

```
interface xe3
qos map-profile dscp-to-queue profile1
!
```

Configuration for DSCP Preserve Global

Topology

The topology above illustrates the DSCP Preserve Global feature, with PE1 and PE2 serving as the edge routers for L3VPN services, and the P node acting as a transit router between them. Traffic received on PE1 with a user-defined DSCP value, which can be preserved or modified depending on the configuration applied to the ingress PE router.

Figure 101. DSCP Preserve Global

Configure DSCP Preserve and L3VPN Services on PE1 Router

1. Enable QoS Globally: This command enables global QoS functionality, allowing configuration of any QoS-related settings. It is mandatory to enable this feature for any QoS configurations to be effective.

```
PE1(config)#qos enable
PE1(config)#commit
```

2. Configure DSCP Preserve Globally: This command enables preserving the customer's DSCP values across the network.

```
PE1(config)#mpls lsp-encap-dscp-preserve
```

3. Configure QoS profiles or class-map profiles for remarking the incoming customer defined DSCP value to customized DSCP value.

- Method 1: QoS Profile (dscp-to-queue): Use a QoS profile to remap the incoming DSCP value to a user-defined queue and DSCP value.

```
PE1(config)#qos profile dscp-to-queue dscp_profile1
```



```

PE1(config-ingress-dscp-map)#dscp 24 queue 3 dscp 32
PE1(config-ingress-dscp-map)#commit
PE1(config-ingress-dscp-map)#exit
PE1(config)#Interface xe0.100
PE1(config-if)#qos map-profile dscp-to-queue dscp_profile1
PE1(config-if)#commit
PE1(config-if)#exit

```

- **Method 2: Class-Map and Policy-Map:** Use class-map and policy-map to match the incoming DSCP value and set a new DSCP value.

```

PE1(config)#class-map type qos match-any class1
PE1(config-qos-match-any)#match dscp 24
PE1(config-qos-match-any)#commit
PE1(config-qos-match-any)#exit

PE1(config)#policy-map type qos pmap1
PE1(config-pmap-qos)#class type qos class1
PE1(config-pmap-qos)#set dscp 32
PE1(config-pmap-qos)#commit
PE1(config-pmap-qos)#exit

PE1(config)#Interface xe0.100
PE1(config-if)#service-policy type qos input pmap1
PE1(config-if)#commit

```

4. Configure EXP Encapsulation on the Egress Interface: Modify the incoming EXP value to a user-defined EXP value for MPLS traffic.

- **Example config on Qumran 1:**

```

PE1(config)#qos profile exp-encap exp_profile1
PE1(config-egress-exp-encap-map)#13 dscp 32 exp 5
PE1(config-egress-exp-encap-map)#commit
PE1(config-egress-exp-encap-map)#exit
PE1(config)#interface ce46/3
PE1(config-if)#label-switching
PE1(config-if)#qos map-profile exp-encap exp_profile1
PE1(config-if)#commit
PE1(config-if)#exit

```

- **Example config on Qumran 2:**

```

PE1(config)#qos profile queue-to-exp exp_profile1
PE1(config-egress-queue-exp-map)#queue 3 color all exp 5
PE1(config-egress-queue-exp-map)#commit
PE1(config-egress-queue-exp-map)#exit
PE1(config)#interface ce46/3
PE1(config-if)#label-switching
PE1(config-if)#qos map-profile queue-to-exp exp_profile1

```

5. Define an IP VRF: Create a VRF with Route Distinguisher (RD) and Route Target (RT) to isolate customer routing information.

```

PE1(config)# ip vrf vrf100
PE1(config-vrf)# rd 1.1.1.1:100
PE1(config-vrf)# route-target both 100:100
PE1(config-vrf)#commit

```

6. Associate Interfaces with VRF: Bind interfaces to the VRF to ensure that the VRF receives the necessary routes.

```

PE1(config)#interface xe0
PE1(config-if)#ip vrf forwarding vrf100
PE1(config-if)#ip address 100.100.100.1/24
PE1(config-if)#exit
PE1(config)#commit

```

7. Configure CE Neighbor for VPN with BGP: Establish a BGP session between the PE router and the CE device, and configure it to advertise routes within the VRF.

```

PE1(config)#router bgp 100
PE1(config-router)# bgp router-id 1.1.1.1
PE1(config-router)#neighbor 3.3.3.3 remote-as 100
PE1(config-router)#neighbor 3.3.3.3 update-source lo
PE1(config-router)#neighbor 3.3.3.3 advertisement-interval 0
PE1(config-router)#address-family ipv4 unicast
PE1(config-router-af)# network 1.1.1.1/32
PE1(config-router-af)#exit-address-family

PE1(config-router)#address-family vpnv4 unicast
PE1(config-router-af)# neighbor 3.3.3.3 activate
PE1(config-router-af)#exit-address-family

PE1(config-router)#address-family ipv4 vrf vrf100
PE1(config-router-af)#redistribute connected
PE1(config-router-af)#neighbor 100.100.100.2 remote-as 100
PE1(config-router-af)#neighbor 100.100.100.2 activate
PE1(config-router-af)#exit-address-family
PE1(config-router)#commit

```

Running Configuration on PE1 Router is as follows:

```

PE1#show running-config
!
! Software version: UFI_S9600-56DX-OcNOS-SP-PLUS-6.5.3.80-Alpha 10/07/2024 14:38:44
!
! Last configuration change at 02:55:39 UTC Tue Oct 15 2024 by root
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile filter qos-ext enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm disable
hardware-profile port-config mode3
!
qos enable
!
qos profile dscp-to-queue dscp_profile1
  dscp 24 queue 3 dscp 32
qos profile queue-to-exp exp_profile1
  queue 3 color all exp 5
!
mpls lsp-encap-dscp-preserve
!
hostname PE1
port ce46 breakout 4X10g
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay

```

```
feature ntp vrf management
ntp enable vrf management
!
ip vrf management
!
ip vrf vrf100
  rd 1.1.1.1:100
  route-target both 100:100
!
router ldp
  targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
  transport-address ipv4 1.1.1.1
!
router rsvp
!
interface ce46/3
  load-interval 30
  ip address 10.10.10.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface lo
  ip address 127.0.0.1/8
  ip address 1.1.1.1/32 secondary
  ipv6 address ::1/128
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
exit
!
router ospf 100
  ospf router-id 1.1.1.1
  network 1.1.1.1/32 area 0.0.0.0
  network 10.10.10.0/24 area 0.0.0.0
!
router bgp 100
  bgp router-id 1.1.1.1
  neighbor 3.3.3.3 remote-as 100
  neighbor 3.3.3.3 update-source lo
  neighbor 3.3.3.3 advertisement-interval 0
!
  address-family ipv4 unicast
    network 1.1.1.1/32
  exit-address-family
!
  address-family vpnv4 unicast
    neighbor 3.3.3.3 activate
  exit-address-family
!
  address-family ipv4 vrf vrf100
    redistribute connected
    neighbor 100.100.100.2 remote-as 100
    neighbor 100.100.100.2 activate
  exit-address-family
!
exit
!
line console 0
```

```
exec-timeout 0
!  
!  
end  
  
PE1#
```

Running Configuration on P Router is as follows:

```
P#show running-config  
!  
! Software version: EC_AS5912-54X-OcNOS-SP-MPLS-6.6.0.104-Alpha 10/13/2024 21:38:49  
!  
! Last configuration change at 02:00:52 UTC Tue Oct 15 2024 by root  
!  
feature netconf-ssh vrf management  
feature netconf-tls vrf management  
no feature netconf-ssh  
no feature netconf-tls  
service password-encryption  
!  
snmp-server enable traps link linkDown  
snmp-server enable traps link linkUp  
!  
hardware-profile statistics ingress-acl enable  
!  
qos enable  
!  
hostname P  
no ip domain-lookup  
ip domain-lookup vrf management  
tfo Disable  
errdisable cause stp-bpdu-guard  
no feature telnet vrf management  
no feature telnet  
feature ssh vrf management  
no feature ssh  
feature dns relay  
ip dns relay  
ipv6 dns relay  
feature ntp vrf management  
ntp enable vrf management  
!  
ip vrf management  
!  
router ldp  
transport-address ipv4 2.2.2.2  
!  
router rsvp  
!  
interface eth0  
ip vrf forwarding management  
ip address dhcp  
!  
interface lo  
ip address 127.0.0.1/8  
ip address 2.2.2.2/32 secondary  
ipv6 address ::1/128  
!  
interface lo.management  
ip vrf forwarding management  
ip address 127.0.0.1/8  
ipv6 address ::1/128  
!  
interface xel  
load-interval 30
```

```
ip address 10.10.10.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface xe9
load-interval 30
ip address 20.20.20.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
exit
!
router ospf 100
ospf router-id 2.2.2.2
network 2.2.2.2/32 area 0.0.0.0
network 10.10.10.0/24 area 0.0.0.0
network 20.20.20.0/24 area 0.0.0.0
!
line console 0
exec-timeout 0
!
!
end

P#
```

Running Configuration on PE2 Router is as follows:

```
PE2#show running-config
!
! Software version: EC_AS7316-26XB-OcNOS-CSR-6.5.3.81-Alpha 10/08/2024 14:39:39
!
! Last configuration change at 07:33:31 UTC Tue Oct 15 2024 by root
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics ingress-acl enable
!
qos enable
!
hostname PE2
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
!
```

```
ip vrf management
!
ip vrf vrf100
  rd 3.3.3.3:100
  route-target both 100:100
!
router ldp
!
router rsvp
!
interface ce0
!
interface ce1
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface lo
  ip address 127.0.0.1/8
  ip address 3.3.3.3/32 secondary
  ipv6 address ::1/128
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
interface xe4.100
  encapsulation dot1q 100
  load-interval 30
  ip vrf forwarding vrf100
  ip address 200.200.200.1/24
  mtu 9216
!
interface xe9
  load-interval 30
  ip address 20.20.20.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
!
exit
!
router ospf 100
  ospf router-id 3.3.3.3
  network 3.3.3.3/32 area 0.0.0.0
  network 20.20.20.0/24 area 0.0.0.0
!
router bgp 100
  bgp router-id 3.3.3.3
  neighbor 1.1.1.1 remote-as 100
  neighbor 1.1.1.1 update-source lo
  neighbor 1.1.1.1 advertisement-interval 0
!
  address-family ipv4 unicast
    network 3.3.3.3/32
  exit-address-family
!
  address-family vpnv4 unicast
    neighbor 1.1.1.1 activate
  exit-address-family
!
  address-family ipv4 vrf vrf100
    redistribute connected
    neighbor 200.200.200.2 remote-as 100
    neighbor 200.200.200.2 activate
  exit-address-family
```

```

!
exit
!
line console 0
  exec-timeout 0
!
!
end
PE2#

```

Validation

PE1

Below show command gives us the output of ospf neighborhood with transit node.

```

PE1#show ip ospf neighbor
Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID      Pri   State           Dead Time   Address        Interface       Instance ID
2.2.2.2          1    Full/DR         00:00:37    10.10.10.2     ce46/3          0
PE1#

```

```

PE1#sh ip vrf
VRF management, VRF ID: 1, FIB ID 1
MPLS DSCP Preserve Enabled (global)
  Router ID: 10.12.96.29 (automatic)
Interfaces:
  eth0
  lo.management
!
VRF vrf100, VRF ID: 2, FIB ID 2
MPLS DSCP Preserve Enabled (global)
  Router ID: 100.100.100.1 (automatic)
Interfaces:
  lo.vrf100
  xe0.100
!
Total Number of configured IP VRF's: 2
Total Number of all VRF's: 3

```

```

Name              Default RD
management        not set
vrf100            1.1.1.1:100
PE1#

```

```

PE1#show ip bgp summary
BGP router identifier 100.100.100.1, local AS number 100
BGP VRF vrf100 Route Distinguisher: 1.1.1.1:100
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|---------------|---|-----|--------|--------|--------|-----|------|----------|--------------|------|
| 100.100.100.2 | 4 | 100 | 116 | 131 | 5 | 0 | 0 | 00:01:06 | 10 | |

Total number of neighbors 1

```

Total number of Established sessions 1
PE1#

```

```

PE1#show ip bgp vpnv4 all summary

```

```
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|---|-----|--------|--------|--------|-----|------|----------|--------------|------|
| 3.3.3.3 | 4 | 100 | 139 | 141 | 2 | 0 | 0 | 00:00:28 | 1 | |

Total number of neighbors 1

Total number of Established sessions 1

```
BGP VRF vrf100 Route Distinguisher: 1.1.1.1:100
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries
```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|---------------|---|-----|--------|--------|--------|-----|------|----------|--------------|------|
| 100.100.100.2 | 4 | 100 | 116 | 131 | 5 | 0 | 0 | 00:01:12 | 10 | |

Total number of neighbors 1

Total number of Established sessions 1

PE1#

Below show commands gives us the traffic statistics. Xe0.100 is ingress interface and ce46/3 is the egress interface.

PE1#show interface counters rate mbps

| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
|-----------|---------|--------|---------|--------|
| ce46/3 | 0.00 | 0 | 10.36 | 9810 |
| xe0 | 10.21 | 9969 | 0.00 | 0 |
| xe0.100 | 10.25 | 10007 | 0.00 | 0 |

PE1#

Below show command gives us the queue-stats on the egress interface of the ingress node. From the configuration made, the traffic is expected to flow in queue 3 as shown below.

PE1#show interface ce46/3 counters queue-stats

E - Egress, I - Ingress, Q-Size is in bytes

| Queue/Class-map | Q-Size | Tx pkts | Tx bytes | Dropped pkts | Dropped bytes |
|-----------------|----------------------|-----------|----------|--------------|---------------|
| q0 | (E) 12499968 0 | 0 | 0 | 0 | 0 |
| q1 | (E) 12499968 0 | 0 | 0 | 0 | 0 |
| q2 | (E) 12499968 0 | 0 | 0 | 0 | 0 |
| q3 | (E) 12499968 7931421 | 872455900 | 0 | 0 | 0 |
| q4 | (E) 12499968 0 | 0 | 0 | 0 | 0 |
| q5 | (E) 12499968 0 | 0 | 0 | 0 | 0 |
| q6 | (E) 12499968 377 | 30275 | 0 | 0 | 0 |
| q7 | (E) 12499968 252 | 14180 | 0 | 0 | 0 |

PE1#

Once the end to end traffic is verified, capture the packet on egress interface of the ingress node. Below is the example packet capture.

Below snapshot is the decoded output of the packet capture. Here we can see that EXP bit is 4 and preserved DSCP value is 24.


```

PE1#show qos-profile type dscp-to-queue dscp_profile1
profile name: dscp_profile1
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
  dscp 24 queue 3 dscp 32
Detailed mapping:

```

| INPUT | | | | OUTPUT | | | | INPUT | | | | OUTPUT | | | | INPUT | | | |
|--------|--------|-------|--------|--------|-------|--------|--------|--------|--------|--------|--------|--------|-------|--------|--------|--------|--------|-------|--------|
| OUTPUT | | | | INPUT | | | | OUTPUT | | | | INPUT | | | | OUTPUT | | | |
| DSCP | TC | Color | remark | DSCP | TC | Color | remark | DSCP | TC | Color | remark | DSCP | TC | Color | remark | DSCP | TC | Color | remark |
| Color | remark | DSCP | DSCP | TC | Color | remark | DSCP | Color | remark | DSCP | DSCP | TC | Color | remark | DSCP | Color | remark | DSCP | DSCP |
| (TC) | | | | (TC) | (TC) | | | (TC) | | | | (TC) | | | | (TC) | | | |
| 0 | 0 | green | 0 (0) | 16 | 2 | green | 16 | 16 | 2 | green | 16 | 32 | 4 | green | 32 (4) | 48 | 6 | green | 48 (6) |
| (2) | 1 | 0 | green | 1 (0) | 17 | 2 | green | 17 | 2 | green | 17 | 33 | 4 | green | 33 (4) | 49 | 6 | green | 49 (6) |
| (2) | 2 | 0 | green | 2 (0) | 18 | 2 | green | 18 | 2 | green | 18 | 34 | 4 | green | 34 (4) | 50 | 6 | green | 50 (6) |
| (2) | 3 | 0 | green | 3 (0) | 19 | 2 | green | 19 | 2 | green | 19 | 35 | 4 | green | 35 (4) | 51 | 6 | green | 51 (6) |
| (2) | 4 | 0 | green | 4 (0) | 20 | 2 | yellow | 20 | 2 | yellow | 20 | 36 | 4 | yellow | 36 (4) | 52 | 6 | green | 52 (6) |
| (2) | 5 | 0 | green | 5 (0) | 21 | 2 | green | 21 | 2 | green | 21 | 37 | 4 | green | 37 (4) | 53 | 6 | green | 53 (6) |
| (2) | 6 | 0 | green | 6 (0) | 22 | 2 | yellow | 22 | 2 | yellow | 22 | 38 | 4 | yellow | 38 (4) | 54 | 6 | green | 54 (6) |
| (2) | 7 | 0 | green | 7 (0) | 23 | 2 | green | 23 | 2 | green | 23 | 39 | 4 | green | 39 (4) | 55 | 6 | green | 55 (6) |
| (2) | 8 | 1 | green | 8 (1) | 24 | 3 | green | 32 | 3 | green | 32 | 40 | 5 | green | 40 (5) | 56 | 7 | green | 56 (7) |
| (4) | 9 | 1 | green | 9 (1) | 25 | 3 | green | 25 | 3 | green | 25 | 41 | 5 | green | 41 (5) | 57 | 7 | green | 57 (7) |
| (3) | 10 | 1 | green | 10 (1) | 26 | 3 | green | 26 | 3 | green | 26 | 42 | 5 | green | 42 (5) | 58 | 7 | green | 58 (7) |
| (3) | 11 | 1 | green | 11 (1) | 27 | 3 | green | 27 | 3 | green | 27 | 43 | 5 | green | 43 (5) | 59 | 7 | green | 59 (7) |
| (3) | 12 | 1 | yellow | 12 (1) | 28 | 3 | yellow | 28 | 3 | yellow | 28 | 44 | 5 | green | 44 (5) | 60 | 7 | green | 60 (7) |
| (3) | 13 | 1 | green | 13 (1) | 29 | 3 | green | 29 | 3 | green | 29 | 45 | 5 | green | 45 (5) | 61 | 7 | green | 61 (7) |
| (3) | 14 | 1 | yellow | 14 (1) | 30 | 3 | yellow | 30 | 3 | yellow | 30 | 46 | 5 | green | 46 (5) | 62 | 7 | green | 62 (7) |
| (3) | 15 | 1 | green | 15 (1) | 31 | 3 | green | 31 | 3 | green | 31 | 47 | 5 | green | 47 (5) | 63 | 7 | green | 63 (7) |
| (3) | | | | | | | | | | | | | | | | | | | |

PE1#

```

PE1#show qos-profile interface ce46/3
profile name: default
profile type: dscp-to-queue (Ingress)
mapping:

```

| INPUT | | | | OUTPUT | | | | INPUT | | | | OUTPUT | | | | INPUT | | | |
|--------|--------|-------|--------|--------|-------|--------|--------|--------|--------|-------|--------|--------|-------|--------|--------|--------|--------|-------|--------|
| OUTPUT | | | | INPUT | | | | OUTPUT | | | | INPUT | | | | OUTPUT | | | |
| DSCP | TC | Color | remark | DSCP | TC | Color | remark | DSCP | TC | Color | remark | DSCP | TC | Color | remark | DSCP | TC | Color | remark |
| Color | remark | DSCP | DSCP | TC | Color | remark | DSCP | Color | remark | DSCP | DSCP | TC | Color | remark | DSCP | Color | remark | DSCP | DSCP |
| (TC) | | | | (TC) | (TC) | | | (TC) | | | | (TC) | | | | (TC) | | | |

| 0 | 0 | green | 0 (0) | | 16 | 2 | green | 16 | |
|-----|---|--------|----------|--------|----|----|---------|--------|--|
| (2) | | 32 | 4 green | 32 (4) | | 48 | 6 green | 48 (6) | |
| 1 | 0 | green | 1 (0) | | 17 | 2 | green | 17 | |
| (2) | | 33 | 4 green | 33 (4) | | 49 | 6 green | 49 (6) | |
| 2 | 0 | green | 2 (0) | | 18 | 2 | green | 18 | |
| (2) | | 34 | 4 green | 34 (4) | | 50 | 6 green | 50 (6) | |
| 3 | 0 | green | 3 (0) | | 19 | 2 | green | 19 | |
| (2) | | 35 | 4 green | 35 (4) | | 51 | 6 green | 51 (6) | |
| 4 | 0 | green | 4 (0) | | 20 | 2 | yellow | 20 | |
| (2) | | 36 | 4 yellow | 36 (4) | | 52 | 6 green | 52 (6) | |
| 5 | 0 | green | 5 (0) | | 21 | 2 | green | 21 | |
| (2) | | 37 | 4 green | 37 (4) | | 53 | 6 green | 53 (6) | |
| 6 | 0 | green | 6 (0) | | 22 | 2 | yellow | 22 | |
| (2) | | 38 | 4 yellow | 38 (4) | | 54 | 6 green | 54 (6) | |
| 7 | 0 | green | 7 (0) | | 23 | 2 | green | 23 | |
| (2) | | 39 | 4 green | 39 (4) | | 55 | 6 green | 55 (6) | |
| 8 | 1 | green | 8 (1) | | 24 | 3 | green | 24 | |
| (3) | | 40 | 5 green | 40 (5) | | 56 | 7 green | 56 (7) | |
| 9 | 1 | green | 9 (1) | | 25 | 3 | green | 25 | |
| (3) | | 41 | 5 green | 41 (5) | | 57 | 7 green | 57 (7) | |
| 10 | 1 | green | 10 (1) | | 26 | 3 | green | 26 | |
| (3) | | 42 | 5 green | 42 (5) | | 58 | 7 green | 58 (7) | |
| 11 | 1 | green | 11 (1) | | 27 | 3 | green | 27 | |
| (3) | | 43 | 5 green | 43 (5) | | 59 | 7 green | 59 (7) | |
| 12 | 1 | yellow | 12 (1) | | 28 | 3 | yellow | 28 | |
| (3) | | 44 | 5 green | 44 (5) | | 60 | 7 green | 60 (7) | |
| 13 | 1 | green | 13 (1) | | 29 | 3 | green | 29 | |
| (3) | | 45 | 5 green | 45 (5) | | 61 | 7 green | 61 (7) | |
| 14 | 1 | yellow | 14 (1) | | 30 | 3 | yellow | 30 | |
| (3) | | 46 | 5 green | 46 (5) | | 62 | 7 green | 62 (7) | |
| 15 | 1 | green | 15 (1) | | 31 | 3 | green | 31 | |
| (3) | | 47 | 5 green | 47 (5) | | 63 | 7 green | 63 (7) | |

```
profile name: default
  profile type: dscp-to-dscp (Egress)
  Status: Inactive
mapping:
```

| -----+----- | | | | -----+----- | | | -----+----- | | | |
|-------------|------|-------|-------------|-------------|------|--------|-------------|-------------|------|-------|
| INPUT | | | OUTPUT | INPUT | | | OUTPUT | INPUT | | |
| -----+----- | | | -----+----- | -----+----- | | | -----+----- | -----+----- | | |
| Remark | DSCP | Color | Out DSCP | Remark | DSCP | Color | Out DSCP | Remark | DSCP | Color |
| Out DSCP | (TC) | | | (TC) | | | | (TC) | | |
| -----+----- | | | -----+----- | -----+----- | | | -----+----- | -----+----- | | |
| 0 | (0) | green | 0 | 0 | (0) | yellow | 0 | 0 | | |
| (0) | red | | 0 | | | | | | | |
| 1 | (0) | green | 1 | 1 | (0) | yellow | 1 | 1 | | |
| (0) | red | | 1 | | | | | | | |
| 2 | (0) | green | 2 | 2 | (0) | yellow | 2 | 2 | | |
| (0) | red | | 2 | | | | | | | |
| 3 | (0) | green | 3 | 3 | (0) | yellow | 3 | 3 | | |
| (0) | red | | 3 | | | | | | | |
| 4 | (0) | green | 4 | 4 | (0) | yellow | 4 | 4 | | |
| (0) | red | | 4 | | | | | | | |
| 5 | (0) | green | 5 | 5 | (0) | yellow | 5 | 5 | | |
| (0) | red | | 5 | | | | | | | |
| 6 | (0) | green | 6 | 6 | (0) | yellow | 6 | 6 | | |
| (0) | red | | 6 | | | | | | | |
| 7 | (0) | green | 7 | 7 | (0) | yellow | 7 | 7 | | |
| (0) | red | | 7 | | | | | | | |
| 8 | (1) | green | 8 | 8 | (1) | yellow | 8 | 8 | | |
| (1) | red | | 8 | | | | | | | |
| 9 | (1) | green | 9 | 9 | (1) | yellow | 9 | 9 | | |
| (1) | red | | 9 | | | | | | | |
| 10 | (1) | green | 10 | 10 | (1) | yellow | 12 | 10 | | |

| | | | | | | |
|--------|-------|----|----|--------|--------|---------|
| (1) | red | 14 | | | | |
| 11 (1) | green | 11 | 11 | 11 (1) | yellow | 11 11 |
| (1) | red | 11 | | | | |
| 12 (1) | green | 12 | 12 | 12 (1) | yellow | 12 12 |
| (1) | red | 14 | | | | |
| 13 (1) | green | 13 | 13 | 13 (1) | yellow | 13 13 |
| (1) | red | 13 | | | | |
| 14 (1) | green | 14 | 14 | 14 (1) | yellow | 14 14 |
| (1) | red | 14 | | | | |
| 15 (1) | green | 15 | 15 | 15 (1) | yellow | 15 15 |
| (1) | red | 15 | | | | |
| 16 (2) | green | 16 | 16 | 16 (2) | yellow | 16 16 |
| (2) | red | 16 | | | | |
| 17 (2) | green | 17 | 17 | 17 (2) | yellow | 17 17 |
| (2) | red | 17 | | | | |
| 18 (2) | green | 18 | 18 | 18 (2) | yellow | 20 18 |
| (2) | red | 22 | | | | |
| 19 (2) | green | 19 | 19 | 19 (2) | yellow | 19 19 |
| (2) | red | 19 | | | | |
| 20 (2) | green | 20 | 20 | 20 (2) | yellow | 20 20 |
| (2) | red | 22 | | | | |
| 21 (2) | green | 21 | 21 | 21 (2) | yellow | 21 21 |
| (2) | red | 21 | | | | |
| 22 (2) | green | 22 | 22 | 22 (2) | yellow | 22 22 |
| (2) | red | 22 | | | | |
| 23 (2) | green | 23 | 23 | 23 (2) | yellow | 23 23 |
| (2) | red | 23 | | | | |
| 24 (3) | green | 24 | 24 | 24 (3) | yellow | 24 24 |
| (3) | red | 24 | | | | |
| 25 (3) | green | 25 | 25 | 25 (3) | yellow | 25 25 |
| (3) | red | 25 | | | | |
| 26 (3) | green | 26 | 26 | 26 (3) | yellow | 28 26 |
| (3) | red | 30 | | | | |
| 27 (3) | green | 27 | 27 | 27 (3) | yellow | 27 27 |
| (3) | red | 27 | | | | |
| 28 (3) | green | 28 | 28 | 28 (3) | yellow | 28 28 |
| (3) | red | 30 | | | | |
| 29 (3) | green | 29 | 29 | 29 (3) | yellow | 29 29 |
| (3) | red | 29 | | | | |
| 30 (3) | green | 30 | 30 | 30 (3) | yellow | 30 30 |
| (3) | red | 30 | | | | |
| 31 (3) | green | 31 | 31 | 31 (3) | yellow | 31 31 |
| (3) | red | 31 | | | | |
| 32 (4) | green | 32 | 32 | 32 (4) | yellow | 32 32 |
| (4) | red | 32 | | | | |
| 33 (4) | green | 33 | 33 | 33 (4) | yellow | 33 33 |
| (4) | red | 33 | | | | |
| 34 (4) | green | 34 | 34 | 34 (4) | yellow | 36 34 |
| (4) | red | 38 | | | | |
| 35 (4) | green | 35 | 35 | 35 (4) | yellow | 35 35 |
| (4) | red | 35 | | | | |
| 36 (4) | green | 36 | 36 | 36 (4) | yellow | 36 36 |
| (4) | red | 38 | | | | |
| 37 (4) | green | 37 | 37 | 37 (4) | yellow | 37 37 |
| (4) | red | 37 | | | | |
| 38 (4) | green | 38 | 38 | 38 (4) | yellow | 38 38 |
| (4) | red | 38 | | | | |
| 39 (4) | green | 39 | 39 | 39 (4) | yellow | 39 39 |
| (4) | red | 39 | | | | |
| 40 (5) | green | 40 | 40 | 40 (5) | yellow | 40 40 |
| (5) | red | 40 | | | | |
| 41 (5) | green | 41 | 41 | 41 (5) | yellow | 41 41 |
| (5) | red | 41 | | | | |
| 42 (5) | green | 42 | 42 | 42 (5) | yellow | 42 42 |
| (5) | red | 42 | | | | |
| 43 (5) | green | 43 | 43 | 43 (5) | yellow | 43 43 |
| (5) | red | 43 | | | | |
| 44 (5) | green | 44 | 44 | 44 (5) | yellow | 44 44 |
| (5) | red | 44 | | | | |
| 45 (5) | green | 45 | 45 | 45 (5) | yellow | 45 45 |
| (5) | red | 45 | | | | |

```

46 (5)      red    green    46      46 | 46 (5)      yellow    46 | 46
47 (5)      red    green    47      47 | 47 (5)      yellow    47 | 47
48 (6)      red    green    48      48 | 48 (6)      yellow    48 | 48
49 (6)      red    green    49      49 | 49 (6)      yellow    49 | 49
50 (6)      red    green    50      50 | 50 (6)      yellow    50 | 50
51 (6)      red    green    51      51 | 51 (6)      yellow    51 | 51
52 (6)      red    green    52      52 | 52 (6)      yellow    52 | 52
53 (6)      red    green    53      53 | 53 (6)      yellow    53 | 53
54 (6)      red    green    54      54 | 54 (6)      yellow    54 | 54
55 (6)      red    green    55      55 | 55 (6)      yellow    55 | 55
56 (7)      red    green    56      56 | 56 (7)      yellow    56 | 56
57 (7)      red    green    57      57 | 57 (7)      yellow    57 | 57
58 (7)      red    green    58      58 | 58 (7)      yellow    58 | 58
59 (7)      red    green    59      59 | 59 (7)      yellow    59 | 59
60 (7)      red    green    60      60 | 60 (7)      yellow    60 | 60
61 (7)      red    green    61      61 | 61 (7)      yellow    61 | 61
62 (7)      red    green    62      62 | 62 (7)      yellow    62 | 62
63 (7)      red    green    63      63 | 63 (7)      yellow    63 | 63

PE1#

```

Commands

The DSCP preserve introduces the following configuration commands.

mpls lsp-encap-dscp-preserve

Use this command to preserve DSCP for IP packets encapsulated into MPLS headers when dscp is remarked on access interface. By default, DSCP is not preserved for IP packets encapsulated into MPLS headers.

Use the no parameter with this to unconfigure DSCP preserve.



Note: In Qumran1 devices, DSCP preserve feature is not supported when the transport is Segment Routing.

Command Syntax

```

mpls lsp-encap-dscp-preserve
no mpls lsp-encap-dscp-preserve

```

Default

By default, mpls local packet handling is disabled

Parameters

None

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.5.2 for Qumran 1. Extended the support for Qumran 2 in OcNOS version 6.5.3.

Examples

```
#configure terminal
(config)#mpls lsp-encap-dscp-preserve
(config)#commit

#configure terminal
(config)#no mpls lsp-encap-dscp-preserve
(config)#commit
```


DSCP Preserve Per-VRF

This method of configuration allows Network operator to enable or disable DSCP preserve for specific L3 VPN services in an edge router.

DSCP Preserve Per-VRF Characteristics

The MPLS cloud connects edge routers that provision L3 VPN services for multiple customers. Each edge router can support multiple VPN services, with different QoS (DSCP) remarking needs. Some customers may want to retain the DSCP field in the IP header, while others may need to update it as traffic enters the MPLS network. By default, the DSCP field in customer packets is not preserved when remarking occurs.

Figure 102. DSCP Preserve Topology



Prerequisites

The L3VPN service will operate seamlessly from end to end without any packet loss, using either OSPF or IS-IS as the IGP, with traffic profiles configured based on user-defined DSCP values.

Global level configuration for PE1:

```
config t
qos enable
!
qos profile dscp-to-queue profile1
dscp 24 queue 3 dscp 32
```

interface level configuration for PE1:


```
interface xe3
qos map-profile dscp-to-queue profile1
!
```

Configuration for DSCP Preserve Per-VRF

Topology

The topology illustrates the DSCP preserve feature, with PE1 and PE2 as the edge routers running L3VPN services, and the P router functioning as a transit router between them. PE1 receiving L3VPN and 6PE traffic with user-defined DSCP values, which can be either preserved or modified depending on the configurations applied at the ingress PE router. Multiple VRFs will be configured, each representing a different customer, where DSCP-preserve can be enabled or disabled at the per-VRF level.

Figure 103. Per- VRF DSCP Preserve



Configure DSCP Preserve and L3VPN Services on PE1 Router

1. Enable QoS Globally: This command is mandatory to configure any QoS-related settings. Without it, QoS configurations will not function.

```
PE1(config)# qos enable
PE1(config)# commit
```

2. (Optional) Configure DSCP Preserve globally for L3VPN and 6PE Traffic: This command provides control over remarking user-defined DSCP values. Enabling this will preserve the DSCP value throughout the network. If DSCP preserve to be enabled only for specific L3 VPN services, then this command need not to be configured.

```
PE1(config)#mpls lsp-encap-dscp-preserve
```

3. Configure DSCP Preserve globally for 6PE Traffic: This command provides control over remarking user-defined DSCP values. Enabling this will preserve the DSCP value throughout the network.

```
PE1(config)# mpls 6pe lsp-encap-dscp-preserve (enable | disable)
```

4. Configure QoS Profiles or Class-Map Profiles for remarking incoming DSCP values as follows:

- Method 1: QoS Profile: Use the “dscp-to-queue” profile to remark the incoming DSCP value via a user-defined queue.

```
PE1(config)# qos profile dscp-to-queue dscp_profile1
PE1(config-ingress-dscp-map)# dscp 24 queue 3 dscp 32
PE1(config-ingress-dscp-map)# commit
PE1(config-ingress-dscp-map)# exit
PE1(config)# interface xe0.100
PE1(config-if)# qos map-profile dscp-to-queue dscp_profile1
PE1(config-if)# commit
PE1(config-if)# exit
```

- Method 2: Class-Map/Policy-Map: Use the “match” criteria to remark the incoming DSCP value to a user-defined DSCP value.

```

PE1(config)# class-map type qos match-any class1
PE1(config-qos-match-any)# match dscp 24
PE1(config-qos-match-any)# commit
PE1(config-qos-match-any)# exit
PE1(config)# policy-map type qos pmap1
PE1(config-pmap-qos)# class type qos class1
PE1(config-pmap-qos)# set dscp 32
PE1(config-pmap-qos)# commit
PE1(config-pmap-qos)# exit
PE1(config)# interface xe0.100
PE1(config-if)# service-policy type qos input pmap1
PE1(config-if)# commit

```

5. Configure EXP Encapsulation on the Egress Interface: This profile modifies the incoming EXP value to the user-defined EXP value, which is then reflected in the MPLS EXP bit.

- Example Config on Qumran 1:

```

PE1(config)# qos profile exp-encap exp_profile1
PE1(config-egress-exp-encap-map)# 13 dscp 32 exp 5
PE1(config-egress-exp-encap-map)# commit
PE1(config-egress-exp-encap-map)# exit
PE1(config)# interface ce46/3
PE1(config-if)# label-switching
PE1(config-if)# qos map-profile exp-encap exp_profile1
PE1(config-if)# commit
PE1(config-if)# exit

```

- Example Config on Qumran 2:

```

PE1(config)# qos profile queue-to-exp exp_profile1
PE1(config-egress-queue-exp-map)# queue 3 color all exp 5
PE1(config-egress-queue-exp-map)# commit
PE1(config-egress-queue-exp-map)# exit
PE1(config)# interface ce46/3
PE1(config-if)# label-switching
PE1(config-if)# qos map-profile queue-to-exp exp_profile1

```

6. Define Multiple IP VRFs: Configure VRFs with appropriate Route Distinguisher (RD) and Route Target (RT) values to isolate IP address routing for multiple customers.

```

PE1(config)# ip vrf vrf100
PE1(config-vrf)# rd 1.1.1.1:100
PE1(config-vrf)# route-target both 100:100
PE1(config-vrf)# commit

PE1(config)# ip vrf vrf101
PE1(config-vrf)# rd 1.1.1.1:101
PE1(config-vrf)# route-target both 101:101
PE1(config-vrf)# commit

PE1(config)# ip vrf vrf102
PE1(config-vrf)# rd 1.1.1.1:102
PE1(config-vrf)# route-target both 102:102
PE1(config-vrf)# commit

```

7. Associate Interfaces to VRFs: After defining the VRFs, assign interfaces to them to allow the PE router to recognize which interfaces belong to each VRF.

```

PE1(config)# interface xe0.100
PE1(config-if)# encapsulation dot1q 100
PE1(config-if)# ip vrf forwarding vrf100
PE1(config-if)# ip address 100.100.100.1/24
PE1(config-if)# exit
PE1(config)# commit

PE1(config)# interface xe0.101
PE1(config-if)# encapsulation dot1q 101

```



```

PE1(config-if)# ip vrf forwarding vrf101
PE1(config-if)# ip address 100.100.101.1/24
PE1(config-if)# exit
PE1(config)# commit

PE1(config)# interface xe0.102
PE1(config-if)# encapsulation dot1q 102
PE1(config-if)# ip vrf forwarding vrf102
PE1(config-if)# ip address 100.100.102.1/24
PE1(config-if)# exit
PE1(config)# commit

PE1(config)# interface xe0.103
PE1(config-if)# encapsulation dot1q 103
PE1(config-if)# ipv6 address 1001::2/64
PE1(config-if)# exit
PE1(config)# commit

```

8. **Configure CE Neighbor for VPN Using BGP:** To provide a VPN service, configure the PE router to associate any routing information learned from a VPN customer interface with a particular VRF.

```

PE1(config)# router bgp 100
PE1(config)# bgp router-id 1.1.1.1
PE1(config)# neighbor 3.3.3.3 remote-as 100
PE1(config)# neighbor 3.3.3.3 update-source lo
PE1(config)# neighbor 3.3.3.3 advertisement-interval 0
PE1(config)# neighbor 1001::2 remote-as 100
PE1(config-router)# address-family ipv4 unicast
PE1(config-router-af)# redistribute connected
PE1(config-router-af)# exit-address-family
PE1(config-router)# address-family ipv4 vrf vrf100
PE1(config-router)# neighbor 100.100.100.2 remote-as 100
PE1(config-router-af)# neighbor 100.100.100.2 activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# address-family ipv4 vrf vrf101
PE1(config-router)# neighbor 100.100.101.2 remote-as 100
PE1(config-router-af)# neighbor 100.100.101.2 activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# address-family ipv4 vrf vrf102
PE1(config-router)# neighbor 100.100.102.2 remote-as 100
PE1(config-router-af)# neighbor 100.100.102.2 activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# address-family ipv6 unicast
PE1(config-router-af)# neighbor 1001::2 activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# commit

```

9. Now enable or disable dscp-preserve based on requirement globally and per-vrf level for L3VPN and 6PE traffic.

```

PE1(config)#ip vrf vrf100
PE1(config-vrf)#mpls lsp-encap-dscp-preserve enable
PE1(config)#ip vrf vrf101
PE1(config-vrf)#mpls lsp-encap-dscp-preserve disable
PE1(config)#ip vrf vrf102
PE1(config-vrf)#mpls lsp-encap-dscp-preserve enable

```

Running Configuration on PE1 Router is as follows:

```

PE1#show running-config
!
! Software version: UFI_S9600-56DX-OcNOS-SP-PLUS-6.5.3.80-Alpha 10/07/2024 14:38:44
!
! Last configuration change at 02:55:39 UTC Tue Oct 15 2024 by root
!
feature netconf-ssh vrf management
feature netconf-tls vrf management

```

```
no feature netconf-ssh
no feature netconf-tls
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile filter qos-ext enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm disable
hardware-profile port-config mode3
!
qos enable
qos profile dscp-to-queue dscp_profile1
  dscp 24 queue 3 dscp 32
qos profile queue-to-exp exp_profile1
  queue 3 color all exp 5
!
mpls lsp-encap-dscp-preserve
!
hostname PE1
port ce46 breakout 4X10g
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
!
ip vrf management
!
ip vrf vrf100
  rd 1.1.1.1:100
  route-target both 100:100
!
router ldp
  targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
  transport-address ipv4 1.1.1.1
!
router rsvp
!
interface ce46/3
  load-interval 30
  ip address 10.10.10.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface lo
  ip address 127.0.0.1/8
  ip address 1.1.1.1/32 secondary
  ipv6 address ::1/128
!
interface lo.management
  ip vrf forwarding management
```

```
ip address 127.0.0.1/8
ipv6 address ::1/128
!
exit
!
router ospf 100
  ospf router-id 1.1.1.1
  network 1.1.1.1/32 area 0.0.0.0
  network 10.10.10.0/24 area 0.0.0.0
!
router bgp 100
  bgp router-id 1.1.1.1
  neighbor 3.3.3.3 remote-as 100
  neighbor 3.3.3.3 update-source lo
  neighbor 3.3.3.3 advertisement-interval 0
!
  address-family ipv4 unicast
    network 1.1.1.1/32
  exit-address-family
!
  address-family vpnv4 unicast
    neighbor 3.3.3.3 activate
  exit-address-family
!
  address-family ipv4 vrf vrf100
    redistribute connected
    neighbor 100.100.100.2 remote-as 100
    neighbor 100.100.100.2 activate
  exit-address-family
!
exit
!
line console 0
  exec-timeout 0
!
!
end

PE1#
```

Running Configuration on P Router is as follows:

```
P#show running-config
!
! Software version: EC_AS5912-54X-OcNOS-SP-MPLS-6.6.0.104-Alpha 10/13/2024 21:38:49
!
! Last configuration change at 02:00:52 UTC Tue Oct 15 2024 by root
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
no feature netconf-ssh
no feature netconf-tls
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics ingress-acl enable
!
qos enable
!
hostname P
no ip domain-lookup
ip domain-lookup vrf management
tfo Disable
errdisable cause stp-bpdu-guard
```

```
no feature telnet vrf management
no feature telnet
feature ssh vrf management
no feature ssh
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
!
ip vrf management
!
router ldp
  transport-address ipv4 2.2.2.2
!
router rsvp
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface lo
  ip address 127.0.0.1/8
  ip address 2.2.2.2/32 secondary
  ipv6 address ::1/128
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
interface xel
  load-interval 30
  ip address 10.10.10.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface xe9
  load-interval 30
  ip address 20.20.20.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
exit
!
router ospf 100
  ospf router-id 2.2.2.2
  network 2.2.2.2/32 area 0.0.0.0
  network 10.10.10.0/24 area 0.0.0.0
  network 20.20.20.0/24 area 0.0.0.0
!
line console 0
  exec-timeout 0
!
!
end

P#
```

Running Configuration on PE2 Router is as follows:

```
PE2#show running-config
```

```
!  
! Software version: EC_AS7316-26XB-OcNOS-CSR-6.5.3.81-Alpha 10/08/2024 14:39:39  
!  
! Last configuration change at 07:33:31 UTC Tue Oct 15 2024 by root  
!  
feature netconf-ssh vrf management  
feature netconf-tls vrf management  
no feature netconf-ssh  
no feature netconf-tls  
service password-encryption  
!  
snmp-server enable traps link linkDown  
snmp-server enable traps link linkUp  
!  
hardware-profile statistics ingress-acl enable  
!  
qos enable  
!  
hostname PE2  
no ip domain-lookup  
ip domain-lookup vrf management  
tfo Disable  
errdisable cause stp-bpdu-guard  
no feature telnet vrf management  
no feature telnet  
feature ssh vrf management  
no feature ssh  
feature dns relay  
ip dns relay  
ipv6 dns relay  
feature ntp vrf management  
ntp enable vrf management  
!  
ip vrf management  
!  
ip vrf vrf100  
  rd 3.3.3.3:100  
  route-target both 100:100  
!  
router ldp  
!  
router rsvp  
!  
interface ce0  
!  
interface ce1  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 3.3.3.3/32 secondary  
  ipv6 address ::1/128  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xe4.100  
  encapsulation dot1q 100  
  load-interval 30  
  ip vrf forwarding vrf100  
  ip address 200.200.200.1/24  
  mtu 9216  
!
```

```

interface xe9
  load-interval 30
  ip address 20.20.20.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
!
exit
!
router ospf 100
  ospf router-id 3.3.3.3
  network 3.3.3.3/32 area 0.0.0.0
  network 20.20.20.0/24 area 0.0.0.0
!
router bgp 100
  bgp router-id 3.3.3.3
  neighbor 1.1.1.1 remote-as 100
  neighbor 1.1.1.1 update-source lo
  neighbor 1.1.1.1 advertisement-interval 0
!
  address-family ipv4 unicast
    network 1.1.1.1/32
  exit-address-family
!
  address-family vpnv4 unicast
    neighbor 1.1.1.1 activate
  exit-address-family
!
  address-family ipv4 vrf vrf100
    redistribute connected
    neighbor 200.200.200.2 remote-as 100
    neighbor 200.200.200.2 activate
  exit-address-family
!
exit
!
line console 0
  exec-timeout 0
!
!
end
PE2#

```

Validation

```

PE1#show ip ospf neighbor
Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID      Pri   State           Dead Time   Address      Interface    Instance ID
2.2.2.2          1    Full/DR         00:00:37   10.10.10.2   ce46/3       0
PE1#

```

```

PE1#sh ip vrf
VRF management, VRF ID: 1, FIB ID 1
MPLS DSCP Preserve Enabled (global)
Router ID: 10.12.96.29 (automatic)
Interfaces:
  eth0
  lo.management
!
VRF vrf100, VRF ID: 2, FIB ID 2
MPLS DSCP Preserve Enabled (global)
Router ID: 100.100.100.1 (automatic)
Interfaces:
  lo.vrf100

```

```

xe0.100
!
Total Number of configured IP VRF's: 2
Total Number of all VRF's: 3

Name                               Default RD
management                         not set
vrf100                             1.1.1.1:100
PE1#

PE1#show ip bgp summary
BGP router identifier 100.100.100.1, local AS number 100
BGP VRF vrf100 Route Distinguisher: 1.1.1.1:100
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries

Neighbor      V    AS    MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd   Desc
100.100.100.2  4    100     116      131      5        0      0   00:01:06         10

Total number of neighbors 1

Total number of Established sessions 1
PE1#

PE1#show ip bgp vpnv4 all summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries

Neighbor      V    AS    MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd   Desc
3.3.3.3        4    100     139      141      2        0      0   00:00:28         1

Total number of neighbors 1

Total number of Established sessions 1

BGP VRF vrf100 Route Distinguisher: 1.1.1.1:100
BGP table version is 5
1 BGP AS-PATH entries
0 BGP community entries

Neighbor      V    AS    MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd   Desc
100.100.100.2  4    100     116      131      5        0      0   00:01:12         10

Total number of neighbors 1

Total number of Established sessions 1
PE1#

PE1#show interface counters rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
| ce46/3    | 0.00    | 0       | 10.36   | 9810    |
| xe0       | 10.21   | 9969    | 0.00    | 0       |
| xe0.100   | 10.25   | 10007   | 0.00    | 0       |
PE1#

PE1#show interface ce46/3 counters queue-stats
E - Egress, I - Ingress, Q-Size is in bytes
+-----+-----+-----+-----+-----+-----+
| Queue/Class-map | Q-Size | Tx pkts | Tx bytes | Dropped pkts | Dropped bytes |
+-----+-----+-----+-----+-----+-----+

```

| | | | | | | |
|----|-----|----------|---------|-----------|---|---|
| q0 | (E) | 12499968 | 0 | 0 | 0 | 0 |
| q1 | (E) | 12499968 | 0 | 0 | 0 | 0 |
| q2 | (E) | 12499968 | 0 | 0 | 0 | 0 |
| q3 | (E) | 12499968 | 7931421 | 872455900 | 0 | 0 |
| q4 | (E) | 12499968 | 0 | 0 | 0 | 0 |
| q5 | (E) | 12499968 | 0 | 0 | 0 | 0 |
| q6 | (E) | 12499968 | 377 | 30275 | 0 | 0 |
| q7 | (E) | 12499968 | 252 | 14180 | 0 | 0 |

PE1#

Once the end to end traffic is verified, capture the packet on egress interface of the ingress node.
Below is the example packet capture.

100.100.100.2 → 200.200.200.2 IPv4 Unknown (253)

| | | | | | | | | | | | | | | | |
|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| 68 | 21 | 5F | 1F | 52 | 21 | 5C | 07 | 58 | 51 | 13 | 42 | 88 | 47 | 06 | 68 |
| 08 | 40 | 06 | B8 | 09 | 40 | 45 | 60 | 00 | 6A | CB | 85 | 00 | 00 | FE | FD |
| 96 | 7F | 64 | 64 | 64 | 02 | C8 | C8 | C8 | 02 | 00 | 00 | 00 | 00 | 00 | 00 |
| 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 |
| 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 |
| 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 |
| 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 00 | 7A | 50 | DE | BE |
| BB | F7 | F9 | 31 | 20 | D9 | 5F | 8C | 19 | 3E | FB | D4 | 91 | B2 | CA | 9B |

4 protocols in packet:

EthernetMPLSMPLSIPv4

Below snapshot is the decoded output of the packet capture. Here we can see that EXP bit is 4 and preserved DSCP value is 24.

- **Frame 1: 128 bytes on wire (1024 bits)**
- **Ethernet II**
- **MultiProtocol Label Switching Header**
 - * 0000 0110 0110 1000 0000 = MPLS Label: 26240 (0x06680)
 - * 100. = MPLS Experimental Bits: 4
 - * 0 = MPLS Bottom Of Label Stack: 0
 - * 0100 0000 = MPLS TTL: 64
- **MultiProtocol Label Switching Header**
- **Internet Protocol Version 4**
 - * 0100 = Version: 4
 - * 0101 = Header Length: 20 bytes (5)
 - * **Differentiated Services Field: 0x60 (DSCP: CS3, ECN: Not-ECT)**
 - * 0110 00.. = Differentiated Services Codepoint: **Class Selector 3 (24)**
 - *00 = Explicit Congestion Notification: Not ECN-Capable Transport (0)
 - * Total Length: 106
 - * Identification: 0xcb85 (52101)
 - * 000. = Flags: 0x0
 - * ...0 0000 0000 0000 = Fragment Offset: 0
 - * Time to Live: 254
 - * Protocol: Unknown (253)
 - * Header Checksum: 0x967f
 - * Header checksum status: Unverified
 - * Source Address: 100.100.100.2
 - * Destination Address: 200.200.200.2
 - * Stream index: 0

```

PE1#show qos-profile type dscp-to-queue dscp_profile1
profile name: dscp_profile1
profile type: dscp-to-queue
profile attached to 1 instances
configured mapping:
  dscp 24 queue 3 dscp 32
Detailed mapping:

```

| INPUT | | | | OUTPUT | | | | INPUT | | | | OUTPUT | | | | INPUT | | | |
|--------|--------|-------|--------|--------|-------|--------|--------|--------|--------|-------|--------|--------|-------|--------|--------|--------|--------|-------|--------|
| OUTPUT | | | | INPUT | | | | OUTPUT | | | | INPUT | | | | OUTPUT | | | |
| DSCP | TC | Color | remark | DSCP | TC | Color | remark | DSCP | TC | Color | remark | DSCP | TC | Color | remark | DSCP | TC | Color | remark |
| Color | remark | DSCP | DSCP | TC | Color | remark | DSCP | Color | remark | DSCP | DSCP | TC | Color | remark | DSCP | Color | remark | DSCP | DSCP |
| (TC) | | | | (TC) | (TC) | | | (TC) | | | | (TC) | | | | (TC) | | | |
| 0 | 0 | green | 0 (0) | 16 | 2 | green | 16 | 0 | 0 | green | 0 (0) | 16 | 2 | green | 16 | 0 | 0 | green | 0 (0) |
| (2) | 32 | 4 | green | 32 (4) | 48 | 6 | green | 48 | (6) | 1 | 0 | green | 1 (0) | 17 | 2 | green | 17 | 1 | 0 |
| (2) | 33 | 4 | green | 33 (4) | 49 | 6 | green | 49 | (6) | 2 | 0 | green | 2 (0) | 18 | 2 | green | 18 | 2 | 0 |
| (2) | 34 | 4 | green | 34 (4) | 50 | 6 | green | 50 | (6) | 3 | 0 | green | 3 (0) | 19 | 2 | green | 19 | 3 | 0 |
| (2) | 35 | 4 | green | 35 (4) | 51 | 6 | green | 51 | (6) | 4 | 0 | green | 4 (0) | 20 | 2 | yellow | 20 | 4 | 0 |
| (2) | 36 | 4 | yellow | 36 (4) | 52 | 6 | green | 52 | (6) | 5 | 0 | green | 5 (0) | 21 | 2 | green | 21 | 5 | 0 |
| (2) | 37 | 4 | green | 37 (4) | 53 | 6 | green | 53 | (6) | 6 | 0 | green | 6 (0) | 22 | 2 | yellow | 22 | 6 | 0 |
| (2) | 38 | 4 | yellow | 38 (4) | 54 | 6 | green | 54 | (6) | | | | | | | | | | |

```

 7      0      green      7 (0) | 23      2      green      23
(2) | 39      4      green      39 (4) | 55      6      green      55 (6)
 8      1      green      8 (1) | 24      3      green      32
(4) | 40      5      green      40 (5) | 56      7      green      56 (7)
 9      1      green      9 (1) | 25      3      green      25
(3) | 41      5      green      41 (5) | 57      7      green      57 (7)
10      1      green      10 (1) | 26      3      green      26
(3) | 42      5      green      42 (5) | 58      7      green      58 (7)
11      1      green      11 (1) | 27      3      green      27
(3) | 43      5      green      43 (5) | 59      7      green      59 (7)
12      1      yellow     12 (1) | 28      3      yellow     28
(3) | 44      5      green      44 (5) | 60      7      green      60 (7)
13      1      green      13 (1) | 29      3      green      29
(3) | 45      5      green      45 (5) | 61      7      green      61 (7)
14      1      yellow     14 (1) | 30      3      yellow     30
(3) | 46      5      green      46 (5) | 62      7      green      62 (7)
15      1      green      15 (1) | 31      3      green      31
(3) | 47      5      green      47 (5) | 63      7      green      63 (7)

```

PE1#

```

PE1#show qos-profile interface ce46/3
profile name: default
profile type: dscp-to-queue (Ingress)
mapping:

```

```

-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| INPUT | | OUTPUT | OUTPUT | INPUT | INPUT | OUTPUT | OUTPUT | INPUT |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+
| DSCP | TC | Color | remark DSCP | DSCP | TC | Color | remark DSCP | DSCP | TC |
Color | remark DSCP | DSCP | (TC) | Color | remark DSCP | DSCP | (TC) |
(TC) | | | | | | | | | |
-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
 0      0      green      0 (0) | 16      2      green      16
(2) | 32      4      green      32 (4) | 48      6      green      48 (6)
 1      0      green      1 (0) | 17      2      green      17
(2) | 33      4      green      33 (4) | 49      6      green      49 (6)
 2      0      green      2 (0) | 18      2      green      18
(2) | 34      4      green      34 (4) | 50      6      green      50 (6)
 3      0      green      3 (0) | 19      2      green      19
(2) | 35      4      green      35 (4) | 51      6      green      51 (6)
 4      0      green      4 (0) | 20      2      yellow     20
(2) | 36      4      yellow     36 (4) | 52      6      green      52 (6)
 5      0      green      5 (0) | 21      2      green      21
(2) | 37      4      green      37 (4) | 53      6      green      53 (6)
 6      0      green      6 (0) | 22      2      yellow     22
(2) | 38      4      yellow     38 (4) | 54      6      green      54 (6)
 7      0      green      7 (0) | 23      2      green      23
(2) | 39      4      green      39 (4) | 55      6      green      55 (6)
 8      1      green      8 (1) | 24      3      green      24
(3) | 40      5      green      40 (5) | 56      7      green      56 (7)
 9      1      green      9 (1) | 25      3      green      25
(3) | 41      5      green      41 (5) | 57      7      green      57 (7)
10      1      green      10 (1) | 26      3      green      26
(3) | 42      5      green      42 (5) | 58      7      green      58 (7)
11      1      green      11 (1) | 27      3      green      27
(3) | 43      5      green      43 (5) | 59      7      green      59 (7)
12      1      yellow     12 (1) | 28      3      yellow     28
(3) | 44      5      green      44 (5) | 60      7      green      60 (7)
13      1      green      13 (1) | 29      3      green      29
(3) | 45      5      green      45 (5) | 61      7      green      61 (7)
14      1      yellow     14 (1) | 30      3      yellow     30
(3) | 46      5      green      46 (5) | 62      7      green      62 (7)
15      1      green      15 (1) | 31      3      green      31

```

```
(3) | 47 5 green 47 (5) | 63 7 green 63 (7)
```

```
profile name: default
profile type: dscp-to-dscp (Egress)
Status: Inactive
mapping:
```

| INPUT | | | OUTPUT | INPUT | | | OUTPUT | INPUT | | |
|-------------|-----|-------|----------|-------------|-----|--------|----------|-------------|-----|-------|
| Remark DSCP | | | Out DSCP | Remark DSCP | | | Out DSCP | Remark DSCP | | |
| (TC) | | | | (TC) | | | | (TC) | | |
| | | | | | | | | | | |
| 0 (0) | red | green | 0 | 0 (0) | red | yellow | 0 | 0 (0) | red | green |
| 1 (0) | red | green | 1 | 1 (0) | red | yellow | 1 | 1 (0) | red | green |
| 2 (0) | red | green | 2 | 2 (0) | red | yellow | 2 | 2 (0) | red | green |
| 3 (0) | red | green | 3 | 3 (0) | red | yellow | 3 | 3 (0) | red | green |
| 4 (0) | red | green | 4 | 4 (0) | red | yellow | 4 | 4 (0) | red | green |
| 5 (0) | red | green | 5 | 5 (0) | red | yellow | 5 | 5 (0) | red | green |
| 6 (0) | red | green | 6 | 6 (0) | red | yellow | 6 | 6 (0) | red | green |
| 7 (0) | red | green | 7 | 7 (0) | red | yellow | 7 | 7 (0) | red | green |
| 8 (1) | red | green | 8 | 8 (1) | red | yellow | 8 | 8 (1) | red | green |
| 9 (1) | red | green | 9 | 9 (1) | red | yellow | 9 | 9 (1) | red | green |
| 10 (1) | red | green | 14 | 10 (1) | red | yellow | 12 | 10 (1) | red | green |
| 11 (1) | red | green | 11 | 11 (1) | red | yellow | 11 | 11 (1) | red | green |
| 12 (1) | red | green | 14 | 12 (1) | red | yellow | 12 | 12 (1) | red | green |
| 13 (1) | red | green | 13 | 13 (1) | red | yellow | 13 | 13 (1) | red | green |
| 14 (1) | red | green | 14 | 14 (1) | red | yellow | 14 | 14 (1) | red | green |
| 15 (1) | red | green | 15 | 15 (1) | red | yellow | 15 | 15 (1) | red | green |
| 16 (2) | red | green | 16 | 16 (2) | red | yellow | 16 | 16 (2) | red | green |
| 17 (2) | red | green | 17 | 17 (2) | red | yellow | 17 | 17 (2) | red | green |
| 18 (2) | red | green | 22 | 18 (2) | red | yellow | 20 | 18 (2) | red | green |
| 19 (2) | red | green | 19 | 19 (2) | red | yellow | 19 | 19 (2) | red | green |
| 20 (2) | red | green | 22 | 20 (2) | red | yellow | 20 | 20 (2) | red | green |
| 21 (2) | red | green | 21 | 21 (2) | red | yellow | 21 | 21 (2) | red | green |
| 22 (2) | red | green | 22 | 22 (2) | red | yellow | 22 | 22 (2) | red | green |
| 23 (2) | red | green | 23 | 23 (2) | red | yellow | 23 | 23 (2) | red | green |
| 24 (3) | red | green | 24 | 24 (3) | red | yellow | 24 | 24 (3) | red | green |
| 25 (3) | red | green | 25 | 25 (3) | red | yellow | 25 | 25 (3) | red | green |
| 26 (3) | red | green | 30 | 26 (3) | red | yellow | 28 | 26 (3) | red | green |

| | | | | | | |
|--------|-----|-------|----|-------------|--------|---------|
| 27 (3) | | green | 27 | 27 27 (3) | yellow | 27 27 |
| (3) | red | | | | | |
| 28 (3) | | green | 28 | 28 28 (3) | yellow | 28 28 |
| (3) | red | | 30 | | | |
| 29 (3) | | green | 29 | 29 29 (3) | yellow | 29 29 |
| (3) | red | | | | | |
| 30 (3) | | green | 30 | 30 30 (3) | yellow | 30 30 |
| (3) | red | | | | | |
| 31 (3) | | green | 31 | 31 31 (3) | yellow | 31 31 |
| (3) | red | | | | | |
| 32 (4) | | green | 32 | 32 32 (4) | yellow | 32 32 |
| (4) | red | | | | | |
| 33 (4) | | green | 33 | 33 33 (4) | yellow | 33 33 |
| (4) | red | | | | | |
| 34 (4) | | green | 34 | 34 34 (4) | yellow | 36 34 |
| (4) | red | | 38 | | | |
| 35 (4) | | green | 35 | 35 35 (4) | yellow | 35 35 |
| (4) | red | | | | | |
| 36 (4) | | green | 36 | 36 36 (4) | yellow | 36 36 |
| (4) | red | | 38 | | | |
| 37 (4) | | green | 37 | 37 37 (4) | yellow | 37 37 |
| (4) | red | | | | | |
| 38 (4) | | green | 38 | 38 38 (4) | yellow | 38 38 |
| (4) | red | | | | | |
| 39 (4) | | green | 39 | 39 39 (4) | yellow | 39 39 |
| (4) | red | | | | | |
| 40 (5) | | green | 40 | 40 40 (5) | yellow | 40 40 |
| (5) | red | | | | | |
| 41 (5) | | green | 41 | 41 41 (5) | yellow | 41 41 |
| (5) | red | | | | | |
| 42 (5) | | green | 42 | 42 42 (5) | yellow | 42 42 |
| (5) | red | | | | | |
| 43 (5) | | green | 43 | 43 43 (5) | yellow | 43 43 |
| (5) | red | | | | | |
| 44 (5) | | green | 44 | 44 44 (5) | yellow | 44 44 |
| (5) | red | | | | | |
| 45 (5) | | green | 45 | 45 45 (5) | yellow | 45 45 |
| (5) | red | | | | | |
| 46 (5) | | green | 46 | 46 46 (5) | yellow | 46 46 |
| (5) | red | | | | | |
| 47 (5) | | green | 47 | 47 47 (5) | yellow | 47 47 |
| (5) | red | | | | | |
| 48 (6) | | green | 48 | 48 48 (6) | yellow | 48 48 |
| (6) | red | | | | | |
| 49 (6) | | green | 49 | 49 49 (6) | yellow | 49 49 |
| (6) | red | | | | | |
| 50 (6) | | green | 50 | 50 50 (6) | yellow | 50 50 |
| (6) | red | | | | | |
| 51 (6) | | green | 51 | 51 51 (6) | yellow | 51 51 |
| (6) | red | | | | | |
| 52 (6) | | green | 52 | 52 52 (6) | yellow | 52 52 |
| (6) | red | | | | | |
| 53 (6) | | green | 53 | 53 53 (6) | yellow | 53 53 |
| (6) | red | | | | | |
| 54 (6) | | green | 54 | 54 54 (6) | yellow | 54 54 |
| (6) | red | | | | | |
| 55 (6) | | green | 55 | 55 55 (6) | yellow | 55 55 |
| (6) | red | | | | | |
| 56 (7) | | green | 56 | 56 56 (7) | yellow | 56 56 |
| (7) | red | | | | | |
| 57 (7) | | green | 57 | 57 57 (7) | yellow | 57 57 |
| (7) | red | | | | | |
| 58 (7) | | green | 58 | 58 58 (7) | yellow | 58 58 |
| (7) | red | | | | | |
| 59 (7) | | green | 59 | 59 59 (7) | yellow | 59 59 |
| (7) | red | | | | | |
| 60 (7) | | green | 60 | 60 60 (7) | yellow | 60 60 |
| (7) | red | | | | | |
| 61 (7) | | green | 61 | 61 61 (7) | yellow | 61 61 |
| (7) | red | | | | | |
| 62 (7) | | green | 62 | 62 62 (7) | yellow | 62 62 |

| | | | | | | | | |
|--------|-------|-------------|--------|---------|--|--|--|--|
| (7) | red | 62 | | | | | | |
| 63 (7) | green | 63 63 (7) | yellow | 63 63 | | | | |
| (7) | red | 63 | | | | | | |
| PE1# | | | | | | | | |

DSCP Commands

The Per-VRF DSCP preserve introduces the following configuration commands.

mpls lsp-encap-dscp-preserve

Use this command to preserve DSCP for IP packets encapsulated into MPLS headers when DSCP is remarked on access interface. By default, DSCP is not preserved for IP packets encapsulated into MPLS headers.

Use the `no` parameter with this to unconfigure DSCP preserve.



Note: In Qumran1 devices, DSCP preserve feature is not supported when the transport is Segment Routing.

Command Syntax

```
mpls lsp-encap-dscp-preserve (enable | disable)
no mpls lsp-encap-dscp-preserve
```

Parameters

enable

Enable the configuration

disable

Disable the configuration

Default

Disable

Command Mode

VRF mode

Applicability

Introduced in OcNOS version 6.5.3.

Example

Example for enabling or disabling DSCP preserve based on requirement per-vrf level:

```
#configure terminal
(config)#ip vrf vrf100
(config-vrf)#mpls lsp-encap-dscp-preserve (enable | disable)
(config-vrf)#commit
(config-vrf)#exit
```

Example for unconfiguring DSCP preserve:

```
#configure terminal
(config)#ip vrf vrf100
(config-vrf)#no mpls lsp-encap-dscp-preserve
(config-vrf)#commit
(config-vrf)#exit
```

mpls 6pe lsp-encap-dscp-preserve

Use this command to preserve DSCP for IP packets encapsulated into MPLS headers when DSCP is remarked on access interface. By default, DSCP is not preserved for IP packets encapsulated into MPLS headers.

Use the `no` parameter with this to unconfigure DSCP preserve.



Note: In Qumran1 devices, DSCP preserve feature is not supported when the transport is Segment Routing.

Command Syntax

```
mpls 6pe lsp-encap-dscp-preserve (enable | disable)
no mpls 6pe lsp-encap-dscp-preserve
```

Parameters

enable

Enable the configuration

disable

Disable the configuration

Default

Disable

Command Mode

Config mode

Applicability

Introduced in OcNOS version 6.5.3.

Example

Example for enabling or disabling DSCP preserve for 6PE services:

```
#configure terminal
(config)#mpls 6pe lsp-encap-dscp-preserve (enable | disable)
(config)#commit
(config)#exit
```

Example for unconfiguring DSCP preserve for 6PE services:

```
#configure terminal
(config)#no mpls 6pe lsp-encap-dscp-preserve
(config)#commit
(config)#exit
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms/Acronym | Description |
|-------------------|---|
| DSCP | Differentiated Services Code Point (DSCP) is a 6-bit value used to classify the priority of Layer-3 packets upon entry into a network. DSCP values range from 0 to 63, 63 being the highest priority, 0 being best-effort traffic. |
| IS-IS | Intermediate System to Intermediate System (IS-IS) An Interior Gateway Protocol (IGP) that floods link state information throughout a network of routers. Each IS-IS router independently builds a database of the network's topology, aggregating the flooded network information. A Routing Information Base (RIB) is calculated from the database by constructing a shortest path tree (SPT). |
| MPLS | Multi-Protocol Label Switching (MPLS) A method for forwarding packets through a network. MPLS operates between the traditional definitions of Layer 2 (L2) and Layer 3 (L3). |
| QoS | Quality of Service (QoS) The ability to guarantee the delivery, control the bandwidth, set priorities for specific network traffic, and provide an appropriate level of security. QoS provides a level of predictability and control beyond the best effort delivery that a device provides by default. |

IPV6 PROVIDER EDGE ROUTERS (6PE) /MPLS VPN CONFIGURATION

| | |
|--|------|
| 6PE Configuration | 2591 |
| Benefits of 6PE | 2591 |
| IPv6 on Provider Edge Routers | 2591 |
| Configuring 6PE | 2592 |
| Validation | 2597 |
| 6VPE Configuration | 2606 |
| Topology | 2606 |
| Configuration | 2606 |
| Validation | 2611 |
| Bridge Virtual Interface (BVI) Over 6vPE | 2621 |
| Overview | 2621 |
| Prerequisites | 2621 |
| Configuration | 2622 |
| Running Configuration | 2625 |
| Validation | 2628 |
| Glossary | 2630 |

6PE Configuration

This chapter explains about IPv6 islands over IPv4 MPLS using IPv6 Provider Edge Routers (6PE). With this technique, IPv6 islands are connected to each other across an IPv4 backbone enabled with MPLS label stacking while MP-BGP is used to announce the IPv6 routes across these MPLS tunnels. This feature can be implemented with label-switched paths (LSPs) using the Label Distribution Protocol (LDP) or Resource Reservation Protocol (RSVP).

This feature offers the following options to the service providers:

- Connect to other IPv6 networks accessible across the MPLS core.
- Provide access to IPv6 services and resources that the service provider provides.
- Provide IPv6 VPN services without going for the complete overhaul of the existing MPLS/IPv4 core.

The 6PE uses the existing IPv4 MPLS core infrastructure for IPv6 transport. It enables IPv6 sites to communicate with each other over an IPv4 MPLS core network using MPLS label switched paths (LSPs). This feature relies heavily on multiprotocol Border Gateway Protocol (BGP) extensions in the IPv4 network configuration on the provider edge (PE) router to exchange IPv6 reachability information (in addition to an MPLS label) for each IPv6 address prefix. Edge routers are configured as dual-stack, running both IPv4 and IPv6, and use the IPv4 mapped IPv6 address for IPv6 prefix reachability exchange.

Benefits of 6PE


6PE offers the following benefits to service providers:

- Minimal operational cost and risk - No impact on existing IPv4 and MPLS services.
- Only provider edge routers require upgrade - A 6PE router can be an existing PE router or a new one dedicated to IPv6 traffic.
- No impact on IPv6 customer edge (CE) routers - The ISP can connect to any CE router running Static, IGP or EGP.
- Production services ready - An ISP can delegate IPv6 prefixes.
- IPv6 introduction into an existing MPLS service - 6PE routers can be added at any time.

IPv6 on Provider Edge Routers

The 6PE is a technique that provides global IPv6 reachability over IPv4 MPLS. It allows one shared routing table for all other devices. 6PE allows IPv6 domains to communicate with one another over the IPv4 without an explicit tunnel setup, requiring only one IPv4 address per IPv6 domain. While implementing 6PE, the provider edge routers are upgraded to support 6PE, while the rest of the core network is not touched (IPv6 unaware).

This implementation requires no re-configuration of core routers because forwarding is based on labels rather than on the IP header itself. This provides a cost-effective strategy for deploying IPv6. The IPv6 reachability information is exchanged by PE routers using multi-protocol Border Gateway Protocol (mp-iBGP) extensions. 6PE relies on mp-iBGP extensions in the IPv4 network configuration on the PE router to exchange IPv6 reachability information in addition to an MPLS label for each IPv6 address prefix to be advertised. PE routers are configured as dual stacks, running both IPv4 and IPv6, and use the IPv4 mapped IPv6 address for IPv6 prefix reachability exchange. The next hop advertised by the PE router for 6PE prefixes is still the IPv4 address that is used for IPv4 L3 VPN routes.



Note: Modifying the next-hop (NH) of 6PE prefixes through route-map is not supported. To set the next-hop of 6PE prefixes to the local address, use the command `neighbor A.B.C.D next-hop-self`.

The following figure illustrates the 6PE topology.

Configuring 6PE


This section shows the configuration of 6PE.

Topology

As shown in [Figure 104](#):

- CE1 and CE2 are customer edge routers
- 6PE1 and 6PE2 are IPv6 Provider Edge routers
- P is the router at the core of the IPv4 MPLS provider network

Figure 104. 6PE Configuration



Configuration

CE1

| | |
|--|--------------------------|
| <pre>#configure terminal</pre> | Enter configure mode |
| <pre>(config)#interface lo</pre> | Enter interface mode |
| <pre>(config-if)#ip address 44.44.44.44/32 secondary</pre> | Assign the IPv4 address |
| <pre>(config-if)#exit</pre> | Exit interface mode |
| <pre>(config)#interface xe1</pre> | Enter interface mode. |
| <pre>(config-if)#ipv6 address 2001::2/64</pre> | Assign the IPv6 address. |

| | |
|---|--|
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router bgp 200</code> | Enter router BGP mode. |
| <code>(config-router)#bgp router-id 44.44.44.44</code> | Assign router ID |
| <code>(config-router)#neighbor 2001::1 remote-as 100</code> | Configure 6PE1 as an eBGP4+ neighbor |
| <code>(config-router)#address-family ipv6 unicast</code> | Enter Address-Family IPv6 unicast mode |
| <code>(config-router-af)#redistribute static</code> | Redistribute static routes |
| <code>(config-router-af)#neighbor 2001::1 activate</code> | Activate the neighbor in the IPv6 address family |
| <code>(config-router-af)#exit</code> | Exit address family |
| <code>(config-router)#ipv6 route 2ffe::/64 xel</code> | Configure IPV6 static route |
| <code>(config-router)#commit</code> | Commit the transaction |

CE2

| | |
|--|---|
| <code>#configure terminal</code> | Enter configure mode |
| <code>(config)#interface lo</code> | Enter interface mode |
| <code>(config-if)#ip address 66.66.66.66/32 secondary</code> | Assign the IPv4 address |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface xel</code> | Enter Interface mode |
| <code>(config-if)#ipv6 address 3002::2/64</code> | Assign IPv6 address |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#router bgp 300</code> | Enter BGP configure mode |
| <code>(config-router)#bgp router-id 66.66.66.66</code> | Assign router ID |
| <code>(config-router)#neighbor 3002::1 remote-as 100</code> | Configure 6PE2 as an eBGP4+ neighbor. |
| <code>(config-router)#address-family ipv6 unicast</code> | Enter Address-Family IPv6 unicast mode |
| <code>(config-router-af)#redistribute static</code> | Redistribute static routes |
| <code>(config-router-af)#neighbor 3002::1 activate</code> | Activate the neighbor in the IPv6 address family. |
| <code>(config-router-af)#exit</code> | Exit address family |
| <code>(config-router)#ipv6 route 3ffe::/64 xel</code> | Configure IPV6 static route |
| <code>(config-router)#commit</code> | Commit the transaction |

PE1

| | |
|--|---|
| <code>#configure terminal</code> | Enter configure mode |
| <code>(config)#interface xel</code> | Enter Interface mode |
| <code>(config-if)#ipv6 address 2001::1/64</code> | Assign IPv6 address |
| <code>(config-if)#exit</code> | Exit interface mode |
| <code>(config)#interface lo</code> | Enter Interface mode |
| <code>(config-if)#ip address 1.1.1.1/32 secondary</code> | Assign the IP address to loopback interface |

| | |
|---|--|
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router ldp mode. |
| (config-router)#router-id 1.1.1.1 | Configure router-id |
| (config-router)#explicit-null | Configure explicit-null. |
| (config-router)#transport-address ipv4 1.1.1.1 | Configure transport address as loopback address |
| (config-router)#targeted-peer ipv4 3.3.3.3 | Configure targeted peer |
| config-router-targeted-peer)#exit | Exit router targeted mode |
| (config-router)#exit | Exit LDP mode |
| (config)#interface xe2 | Enter Interface mode |
| (config-if)#ip address 20.1.1.1/24 | Assign IPv4 address |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable ldp in interface. |
| (config-if)#exit | Exit interface mode |
| (config)#router ospf | Enter router ospf mode. |
| (config-router)#ospf router-id 1.1.1.1 | Configure ospf router id same as loopback ip address. |
| (config-router)#network 1.1.1.1/32 area 0 (config-router)#network 20.1.1.0/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit from router ospf mode. |
| (config)#mpls label mode 6pe per-prefix | Change label mode to per-prefix, default is per VRF |
| (config)#router bgp 100 | Enter BGP Configure mode. |
| (config-router)#bgp router-id 1.1.1.1 | Configure BGP router-id |
| (config-router)#neighbor 3.3.3.3 remote-as 100 | Configure 6PE2 as an iBGP peer. |
| (config-router)#neighbor 3.3.3.3 update-source lo | Update the source as loopback for iBGP peering with the remote 6PE router. |
| (config-router)#neighbor 2001::2 remote-as 200 | Configure CE1 as eBGP peer |
| (config-router)#address-family ipv4 unicast | Enter address family mode |
| (config-router-af)#neighbor 3.3.3.3 activate | Activate neighbor |
| (config-router-af)#exit | Exit address family mode |
| (config-router)#address-family ipv6 labeled-unicast | Enter IPv6 labeled-unicast Address Family mode. |
| (config-router-af)#neighbor 3.3.3.3 activate | Activate the 6PE neighbor |
| (config-router-af)#exit-address-family | Exit IPv6 LU Address Family mode. |
| (config-router)#address-family ipv6 unicast | Enter the IPv6 address family |
| (config-router-af)#neighbor 2001::2 activate | Activate CE inside IPv6 address family |
| (config-router-af)#redistribute connected | Redistribute the connected routes |
| (config-router-af)#exit | Exit IPv6 Address Family mode. |
| (config-router)#commit | Commit the transaction. |

P1

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 2.2.2.2/32 secondary | Assign the IP address to the loopback interface |
| (config-if)#exit | Exit interface mode |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#router-id 2.2.2.2 | Configure router-id |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure the transport address as the loopback address |
| (config-router)#exit | Exit router LDP mode. |
| (config)#router ospf | Enter router OSPF mode. |
| (config-router)#ospf router-id 2.2.2.2 | Configure the OSPF router ID to be the same as a loopback IP address. |
| (config-router)#network 2.2.2.2/32 area 0 (config-router)#network 20.1.1.2/24 area 0 (config-router)#network 30.1.1.1/24 area 0 | Define the network on which OSPF runs and associate the area ID. |
| (config-router)#exit | Exit from router ospf mode. |
| (config)#interface xe2 | Enter Interface mode |
| (config-if)#ip address 30.1.1.1/24 | Assign IPv4 address |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable ldp in interface. |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe1 | Enter Interface mode |
| (config-if)#ip address 20.1.1.2/24 | Assign IPv4 address |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-ldp ipv4 | Enable ldp in interface. |
| (config-if)#commit | Commit the transaction. |

PE2

| | |
|---|---|
| #configure terminal | Enter configure mode |
| (config)#interface xe2 | Enter Interface mode |
| (config-if)#ipv6 address 3002::1/64 | Assign IPv6 address |
| (config-if)#exit | Exit interface mode |
| (config)#interface lo | Enter Interface mode |
| (config-if)#ip address 3.3.3.3/32 secondary | Assign the IP address to the loopback interface |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router LDP mode. |

| | |
|---|--|
| (config-router)#router-id 3.3.3.3 | Configure router-id |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure the transport address as the loopback address |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#explicit-null | Configure explicit-null. |
| (config-router)#exit | Exit LDP mode |
| (config)#interface xel | Enter Interface mode |
| (config-if)#ip address 30.1.1.2/24 | Assign IPv4 address |
| (config-if)#label-switching | Enable label switching in the interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in the interface. |
| (config-if)#exit | Exit interface mode |
| (config)#mpls label mode 6pe per-prefix | Change label mode to per-prefix, the default is per VRF |
| (config)#router bgp 100 | Enter router BGP mode. |
| (config-router)#bgp router-id 3.3.3.3 | Configure the BGP router ID. |
| (config-router)#neighbor 1.1.1.1 remote-as 100 | Configure 6VPE2 as an iBGP peer. |
| (config-router)#neighbor 1.1.1.1 update-source lo | Update the source as loopback for iBGP peering with the remote 6VPE router. |
| (config-router)#address-family ipv4 unicast | Enter address family mode |
| (config-router-af)#neighbor 1.1.1.1 activate | Activate neighbor |
| (config-router-af)#exit | Exit address family mode |
| (config-router)#neighbor 3002::2 remote-as 300 | Configure CE1 as an eBGP peer |
| (config-router)#address-family ipv6 labeled-unicast | Enter IPv6 labeled-unicast Address Family mode. |
| (config-router-af)#neighbor 1.1.1.1 activate | Activate the 6PE neighbor |
| (config-router-af)#exit-address-family | Exit IPv6 LU Address Family mode. |
| (config-router)#address-family ipv6 unicast | Enter the IPv6 address family |
| (config-router-af)#neighbor 3002::2 activate | Activate CE inside the IPv6 address family |
| (config-router-af)#redistribute connected | Redistribute the connected routes |
| (config-router-af)#exit-address-family | Exit IPv6 Address Family mode. |
| (config-router)#exit | Exit Router mode. |
| (config)#router ospf | Enter OSPF router mode |
| (config-router)#network 3.3.3.3/32 area 0 | Enable OSPF with the specified area ID on interfaces with the IP address that matches the specified network address. |
| (config-router)#network 30.1.1.0/24 area 0 | Enable OSPF with the specified area ID on interfaces with the IP address that matches the specified network address |
| (config-router)#commit | Commit the transaction |

Validation

CE1

```
CE1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:10:32
C      2001::/64 via ::, xe1, 00:46:49
S      2ffe::/64 [1/0] via ::, xe1, 00:35:20
B      3002::/64 [20/0] via fe80::5054:ff:fe29:189d, xe1, 00:02:12
B      3ffe::/64 [20/0] via fe80::5054:ff:fe29:189d, xe1, 00:02:36
C      fe80::/64 via ::, xe3, 01:10:32
#

CE1#show ipv6 bgp summary
BGP router identifier 44.44.44.44, local AS number 200
BGP table version is 8
3 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd
2001::1           4    100    80      83       8       0     0    00:01:45      3

Total number of neighbors 1

Total number of Established sessions 1
```

PE1

```
PE1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:17:11
C      2001::/64 via ::, xe1, 00:40:22
B      2ffe::/64 [20/0] via fe80::5054:ff:fe60:f4e5, xe1, 00:02:37
B      3002::/64 [200/0] via ::ffff:3.3.3.3, 00:03:10
B      3ffe::/64 [200/0] via ::ffff:3.3.3.3, 00:01:07
C      fe80::/64 via ::, xe2, 01:17:11

PE1#show ldp session
Peer IP Address  IF Name    My Role    State          KeepAlive    UpTime
2.2.2.2         xe2        Passive    OPERATIONAL    30           00:06:59

PE1#show bgp ipv6
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|----------------|-----------------------------------|--------|--------|--------|-------|
| *> 1 2001::/64 | :: | 0 | 100 | 32768 | ? |
| *> 1 2ffe::/64 | 2001::2 (fe80::5054:ff:fe60:f4e5) | 0 | 100 | 0 | 200 ? |
| *>i 3002::/64 | ::ffff:3.3.3.3 | 0 | 100 | 0 | ? |
| *>i 3ffe::/64 | ::ffff:3.3.3.3 | 0 | 100 | 0 | 300 ? |

Total number of prefixes 4

PE1#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|-------------|---------|----------|-----------|-----|-------------|-----------|------|
| Intf | ELC | NextHop | | | | | | |
| L> | 2.2.2.2/32 | 3 | 3 | - | - | LSP_ | | |
| DEFAULT | 3 | xe2 | No | 20.1.1.2 | | | | |
| L> | 3.3.3.3/32 | 4 | 4 | - | - | LSP_ | | |
| DEFAULT | 24321 | xe2 | No | 20.1.1.2 | | | | |
| L> | 30.1.1.0/24 | 5 | 3 | - | - | LSP_ | | |
| DEFAULT | 3 | xe2 | No | 20.1.1.2 | | | | |
| B> | 3002::/64 | 2 | 2 | 0 | Yes | LSP_DEFAULT | 24960 | - |
| | No 3.3.3.3 | | | | | | | |
| B> | 3ffe::/64 | 1 | 1 | 0 | Yes | LSP_DEFAULT | 24961 | - |
| | No 3.3.3.3 | | | | | | | |

PE1#show ldp session

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 2.2.2.2 | xe2 | Passive | OPERATIONAL | 30 | 00:06:59 |

PE1#show mpls ftn-table

Primary FTN entry with FEC: 2.2.2.2/32, id: 3, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 3, owner: N/A, Stale: NO, out intf: xe2, out label: 3
 Nexthop addr: 20.1.1.2 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 4, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 4
 Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 4, owner: LDP, Stale: NO, out intf: xe2, out label: 24321
 Nexthop addr: 20.1.1.2 cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 30.1.1.0/24, id: 5, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 3, owner: N/A, Stale: NO, out intf: xe2, out label: 3
 Nexthop addr: 20.1.1.2 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 3002::/64, id: 2, row status: Active, Tunnel-Policy: N/A
 Owner: BGP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Transport Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0

```

Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: BGP, Stale: NO, BGP out intf: xe2, transport out intf: xe2, out
label: 24960
Nexthop addr: 3.3.3.3          cross connect ix: 2, op code: Push and Lookup

```

```

Primary FTN entry with FEC: 3ffe::/64, id: 1, row status: Active, Tunnel-Policy: N/A
Owner: BGP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Transport Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: BGP, Stale: NO, BGP out intf: xe2, transport out intf: xe2, out
label: 24961
Nexthop addr: 3.3.3.3          cross connect ix: 1, op code: Push and Lookup

```

```

PE1#show mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM
      K - CLI ILM, T - MPLS-TP, s - Stitched ILM
      S - SNMP, L - LDP, R - RSVP, C - CRLDP
      B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
      O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
      P - SR Policy, U - unknown

```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- | |
|----------|---------------|--------|----------|-----------|---------|------|-----------|
| Intf/VRF | Nexthop | | LSP-Type | | | | |
| B> | 2001::/64 | 3 | 24960 | Nolabel | N/A | N/A | 127.0.0.1 |
| | LSP_DEFAULT | | | | | | |
| B> | 2ffe::/64 | 4 | 24961 | Nolabel | N/A | N/A | 127.0.0.1 |
| | LSP_DEFAULT | | | | | | |

#

```

PE1#show ip bgp summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|----------|---|-----|--------|--------|--------|-----|------|----------|--------------|
| 3.3.3.3 | 4 | 100 | 42 | 43 | 1 | 0 | 0 | 00:08:40 | 0 |

Total number of neighbors 1

Total number of Established sessions 1

```

PE1#show ipv6 bgp summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 5
3 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|----------|---|-----|--------|--------|--------|-----|------|----------|--------------|
| 2001::2 | 4 | 200 | 93 | 98 | 5 | 0 | 0 | 00:08:33 | 1 |

Total number of neighbors 1

Total number of Established sessions 1

```

PE1#show ip bgp neighbors
BGP neighbor is 3.3.3.3, remote AS 100, local AS 100, internal link
BGP version 4, local router ID 1.1.1.1, remote router ID 3.3.3.3
BGP state = Established, up for 00:08:55
Last read 00:00:21, hold time is 90, keepalive interval is 30 seconds
Neighbor capabilities:
Route refresh: advertised and received (old and new)
Address family IPv4 Unicast: advertised and received
Address family IPv6 Labeled Unicast: advertised and received

```

```

Received 42 messages, 0 notifications, 0 in queue
Sent 43 messages, 1 notifications, 0 in queue
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 5 seconds
Update source is lo
For address family: IPv4 Unicast
BGP table version 1, neighbor version 1
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (both)
0 accepted prefixes
0 announced prefixes

For address family: IPv6 Labeled-Unicast
BGP table version 6, neighbor version 6
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (both)
2 accepted prefixes
2 announced prefixes

Connections established 2; dropped 1
Local host: 1.1.1.1, Local port: 34293
Foreign host: 3.3.3.3, Foreign port: 179
Nexthop: 1.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:09:51, due to Administratively Reset (Cease Notification sent)
Notification Error Message: (Cease/Administratively Reset.)

BGP neighbor is 2001::2, remote AS 200, local AS 100, external link
BGP version 4, local router ID 1.1.1.1, remote router ID 44.44.44.44
BGP state = Established, up for 00:08:45
Last read 00:00:16, hold time is 90, keepalive interval is 30 seconds
Neighbor capabilities:
Route refresh: advertised and received (old and new)
Address family IPv6 Unicast: advertised and received
Received 92 messages, 1 notifications, 0 in queue
Sent 97 messages, 1 notifications, 0 in queue
Route refresh request: received 0, sent 0
Minimum time between advertisement runs is 30 seconds
For address family: IPv6 Unicast
BGP table version 5, neighbor version 5
Index 1, Offset 0, Mask 0x2
Community attribute sent to this neighbor (both)
1 accepted prefixes
3 announced prefixes

Connections established 3; dropped 2
Local host: 2001::1, Local port: 179
Foreign host: 2001::2, Foreign port: 40980
Nexthop: 1.1.1.1
Nexthop global: 2001::1
Nexthop local: fe80::5054:ff:fe29:189d
BGP connection: shared network
Last Reset: 00:08:50, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

```

P1

```

P1#show ldp session
Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
3.3.3.3              xe1        Passive    OPERATIONAL 30         00:10:11
1.1.1.1              xe2        Active     OPERATIONAL 30         00:09:21

P1#show mpls forwarding-table

```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
        B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|------------|----------|----------|-----------|-----|----------|-----------|------|
| Intf | ELC | Nextthop | | | | | | |
| L> | 1.1.1.1/32 | 2 | 2 | - | - | LSP_ | | |
| DEFAULT | 0 | xe2 | No | 20.1.1.1 | | | | |
| L> | 3.3.3.3/32 | 1 | 1 | - | - | LSP_ | | |
| DEFAULT | 0 | xe1 | No | 30.1.1.2 | | | | |

```
Pl#show mpls ilm-table
```

```
Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM, T - MPLS-TP, s - Stitched ILM
        S - SNMP, L - LDP, R - RSVP, C - CRLDP
        B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
        O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
        P - SR Policy, U - unknown
```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|--------|----------|-----------|---------|----------|
| Intf/VRF | Nextthop | | LSP-Type | | | |
| L> | 3.3.3.3/32 | 2 | 24321 | 0 | N/A | xe1 |
| | LSP_DEFAULT | | | | | 30.1.1.2 |
| L> | 1.1.1.1/32 | 1 | 24320 | 0 | N/A | xe2 |
| | LSP_DEFAULT | | | | | 20.1.1.1 |

PE2

```
PE2#show ipv6 route
```

```
IPv6 Routing Table
```

```
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
        O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
        E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,
        N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
        v - vrf leaked
```

```
Timers: Uptime
```

```
IP Route Table for VRF "default"
```

```
C      ::1/128 via ::, lo, 01:24:48
B      2001::/64 [200/0] via ::ffff:1.1.1.1, 00:11:08
B      2ffe::/64 [200/0] via ::ffff:1.1.1.1, 00:10:34
C      3002::/64 via ::, xe2, 00:24:41
B      3ffe::/64 [20/0] via fe80::5054:ff:fef6:c35d, xe2, 00:09:07
C      fe80::/64 via ::, xe3, 01:24:48
```

```
PE2#show mpls ilm-table
```

```
Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM, T - MPLS-TP, s - Stitched ILM
        S - SNMP, L - LDP, R - RSVP, C - CRLDP
        B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
        O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
        P - SR Policy, U - unknown
```

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|--------|----------|-----------|---------|-----------|
| Intf/VRF | Nextthop | | LSP-Type | | | |
| B> | 3002::/64 | 3 | 24960 | Nolabel | N/A | N/A |
| | LSP_DEFAULT | | | | | 127.0.0.1 |
| B> | 3ffe::/64 | 4 | 24961 | Nolabel | N/A | N/A |
| | LSP_DEFAULT | | | | | 127.0.0.1 |

```
PE2#show mpls forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
        B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
```

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|------|-----|--------|----------|-----------|-----|----------|-----------|------|
|------|-----|--------|----------|-----------|-----|----------|-----------|------|

```

Intf    ELC    Nexthop
L>      1.1.1.1/32      3      2      -      -      LSP_
DEFAULT 24320      xe1      No      30.1.1.1
L>      2.2.2.2/32      1      1      -      -      LSP_
DEFAULT 3      xe1      No      30.1.1.1
L>      20.1.1.0/24     2      1      -      -      LSP_
DEFAULT 3      xe1      No      30.1.1.1
B>      2001::/64      4      3      0      Yes    LSP_DEFAULT 24960      -
      No      1.1.1.1
B>      2ffe::/64      5      4      0      Yes    LSP_DEFAULT 24961      -
      No      1.1.1.1

```

PE2#show mpls ftn-table

```

Primary FTN entry with FEC: 1.1.1.1/32, id: 3, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 2
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: xe1, out label: 24320
Nexthop addr: 30.1.1.1      cross connect ix: 3, op code: Push

```

```

Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe1, out label: 3
Nexthop addr: 30.1.1.1      cross connect ix: 2, op code: Push

```

```

Primary FTN entry with FEC: 20.1.1.0/24, id: 2, row status: Active, Tunnel-Policy: N/A
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 1
Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe1, out label: 3
Nexthop addr: 30.1.1.1      cross connect ix: 2, op code: Push

```

```

Primary FTN entry with FEC: 2001::/64, id: 4, row status: Active, Tunnel-Policy: N/A
Owner: BGP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Transport Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: BGP, Stale: NO, BGP out intf: xe1, transport out intf: xe1, out
label: 24960
Nexthop addr: 1.1.1.1      cross connect ix: 4, op code: Push and Lookup

```

```

Primary FTN entry with FEC: 2ffe::/64, id: 5, row status: Active, Tunnel-Policy: N/A
Owner: BGP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Transport Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 4
Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 4, owner: BGP, Stale: NO, BGP out intf: xe1, transport out intf: xe1, out
label: 24961
Nexthop addr: 1.1.1.1      cross connect ix: 5, op code: Push and Lookup

```

PE2#show ldp session

| Peer IP Address | IF Name | My Role | State | KeepAlive | UpTime |
|-----------------|---------|---------|-------------|-----------|----------|
| 2.2.2.2 | xe1 | Active | OPERATIONAL | 30 | 00:12:01 |

```

PE2#show bgp ipv6
BGP table version is 5, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal,
               l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|----------------|-----------------------------------|--------|--------|--------|-------|
| *>i 2001::/64 | ::ffff:1.1.1.1 | 0 | 100 | 0 | ? |
| *>i 2ffe::/64 | ::ffff:1.1.1.1 | 0 | 100 | 0 | 200 ? |
| *> l 3002::/64 | :: | 0 | 100 | 32768 | ? |
| *> l 3ffe::/64 | 3002::2 (fe80::5054:ff:fef6:c35d) | 0 | 100 | 0 | 300 ? |

Total number of prefixes 4

```

PE2#show ip bgp neighbors
BGP neighbor is 1.1.1.1, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 3.3.3.3, remote router ID 1.1.1.1
  BGP state = Established, up for 00:11:54
  Last read 00:00:06, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Labeled Unicast: advertised and received
  Received 50 messages, 0 notifications, 0 in queue
  Sent 50 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
  For address family: IPv4 Unicast
    BGP table version 1, neighbor version 1
    Index 1, Offset 0, Mask 0x2
    Community attribute sent to this neighbor (both)
    0 accepted prefixes
    0 announced prefixes

  For address family: IPv6 Labeled-Unicast
    BGP table version 5, neighbor version 5
    Index 1, Offset 0, Mask 0x2
    Community attribute sent to this neighbor (both)
    2 accepted prefixes
    2 announced prefixes

  Connections established 2; dropped 1
  Local host: 3.3.3.3, Local port: 179
  Foreign host: 1.1.1.1, Foreign port: 34293
  Nexthop: 3.3.3.3
  Nexthop global: ::
  Nexthop local: ::
  BGP connection: non shared network
  Last Reset: 00:12:28, due to Administratively Reset (Cease Notification sent)
  Notification Error Message: (Cease/Administratively Reset.)

```

```

BGP neighbor is 3002::2, remote AS 300, local AS 100, external link
  BGP version 4, local router ID 3.3.3.3, remote router ID 66.66.66.66
  BGP state = Established, up for 00:10:17
  Last read 00:00:25, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv6 Unicast: advertised and received
  Received 61 messages, 2 notifications, 0 in queue
  Sent 68 messages, 2 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
  For address family: IPv6 Unicast
    BGP table version 5, neighbor version 5
    Index 1, Offset 0, Mask 0x2
    Community attribute sent to this neighbor (both)

```

```

1 accepted prefixes
3 announced prefixes

Connections established 3; dropped 2
Local host: 3002::1, Local port: 52758
Foreign host: 3002::2, Foreign port: 179
Nexthop: 3.3.3.3
Nexthop global: 3002::1
Nexthop local: fe80::5054:ff:fe2b:8d4f
BGP connection: shared network
Last Reset: 00:10:22, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

PE2#show ip bgp summary
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd
1.1.1.1            4    100     50      52       1      0      0   00:12:06           0

Total number of neighbors 1

Total number of Established sessions 1

PE2#show ipv6 bgp summary
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 5
3 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd
3002::2           4    300     64      70       5      0      0   00:10:31           1

Total number of neighbors 1

Total number of Established sessions 1

```

CE2

```

CE2#show ipv6 bgp summary
BGP router identifier 66.66.66.66, local AS number 300
BGP table version is 9
3 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv   MsgSen  TblVer   InQ   OutQ   Up/Down   State/PfxRcd
3002::1           4    100     70      67       9      0      0   00:11:35           3

Total number of neighbors 1

Total number of Established sessions 1

CE2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:26:48

```



```
B      2001::/64 [20/0] via fe80::5054:ff:fe2b:8d4f, xe2, 00:11:43
B      2ffe::/64 [20/0] via fe80::5054:ff:fe2b:8d4f, xe2, 00:11:43
C      3002::/64 via ::, xe2, 00:24:47
S      3ffe::/64 [1/0] via ::, xe2, 00:24:05
C      fe80::/64 via ::, xe2, 01:26:48
```

6VPE Configuration


This chapter explains how 6VPE (IPv6 on VPN Provider Edge Routers) can interconnect IPv6 islands over an MPLS- enabled IPv4 cloud. 6VPE enables IPv6 sites to communicate with each other over an MPLS/IPv4 core network using MPLS LSPs. The 6VPE routers exchange IPv6 reachability information over the core using Multi-Protocol Border Gateway Protocol (MP-BGP) over IPv4.

Topology

As shown in [Figure 105](#):

- CE1 and CE2 are customer edge routers
- 6VPE1 and 6VPE2 are IPv6 Provider Edge routers
- P is the router at the core of the IPv4 MPLS provider network.

Figure 105. 6VPE Configuration



Configuration

CE1

| | |
|---|------------------------------|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 44.44.44.44/32 secondary | Assign the IPv4 address |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ipv6 address 2001::2/64 | Assign the IPv6 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#ipv6 route 2ffe::/64 xe1 | Advertise IPv6 static route. |

| | |
|--|---|
| (config)#router bgp 200 | Enter BGP router mode. |
| (config-router)#bgp router-id 44.44.44.44 | Configure bgp router-id |
| (config-router)#neighbor 2001::1 remote-as 100 | Configure 6VPE1 as an eBGP4+ neighbor. |
| (config-router)#address-family ipv6 unicast | Enter address-family IPv6 unicast mode. |
| (config-router-af)#neighbor 2001::1 activate | Activate the neighbor in the IPv6 address family. |
| (config-router-af)#redistribute connected | Redistribute the connected route under address family IPv6 unicast. |
| (config-router-af)#redistribute static | Redistribute static routes. |
| (config-router-af)#commit | Commit the transaction. |

CE2

| | |
|---|---|
| #configure terminal | Enter configure mode |
| (config)#interface lo | Enter interface mode |
| (config-if)#ip address 66.66.66.66/32 secondary | Assign the IPv4 address |
| (config-if)#exit | Exit interface mode |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#ipv6 address 3001::2/64 | Assign the IPv6 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#ipv6 route 3ffe::/64 xe1 | Configure IPV6 static route |
| (config)#router bgp 300 | Enter BGP router mode. |
| (config-router)#bgp router-id 66.66.66.66 | Configure BGP router-id |
| (config-router)#neighbor 3001::1 remote-as 100 | Configure 6VPE1 as an eBGP4+ neighbor. |
| (config-router)#address-family ipv6 unicast | Enter address-family IPv6 unicast mode. |
| (config-router-af)#neighbor 3001::1 activate | Activate the neighbor in the IPv6 address family. |
| (config-router-af)#redistribute connected | Redistribute the connected route under address family IPv6 unicast. |
| (config-router-af)#redistribute static | Redistribute static routes. |
| (config-router-af)#commit | Commit the transaction. |

PE1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#ip vrf IPI | Create a new VRF named IPI. |
| (config-vrf)#rd 1:100 | Assign the route distinguisher (RD) value as 1:100. |
| (config-vrf)#route-target both 100:200 | Import routes between route target (RT) ext-communities 100 and 200. |
| (config-vrf)#router-id 77.77.77.77 | Configure router-id for VRF |

| | |
|--|---|
| <code>(config-vrf)#exit</code> | Exit VRF mode. |
| <code>(config)#interface xe1</code> | Enter interface mode. |
| <code>(config-if)#ip vrf forwarding IPI</code> | Bind the interface connected to the CE router with VRF IPI. |
| <code>(config-if)#ipv6 address 2001::1/64</code> | Assign the IPv6 address. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router bgp 100</code> | Enter BGP router mode. |
| <code>(config-router)#bgp router-id 1.1.1.1</code> | Configure BGP router-id |
| <code>(config-router)#neighbor 3.3.3.3 remote-as 100</code> | Configure 6VPE2 as an iBGP peer. |
| <code>(config-router)#neighbor 3.3.3.3 update-source lo</code> | Update the source as loopback for iBGP peering with the remote 6VPE router. |
| <code>(config-router)#address-family ipv4 unicast</code> | Enter address family mode |
| <code>(config-router-af)#neighbor 3.3.3.3 activate</code> | Activate the neighbor |
| <code>(config-router-af)#exit-address-family</code> | Exit address family mode |
| <code>(config-router)#address-family vpnv6 unicast</code> | Enter VPNv6 address family mode. |
| <code>(config-router-af)#neighbor 3.3.3.3 activate</code> | Activate the 6VPE neighbor so that it can accept VPN IPv6 routes. |
| <code>(config-router-af)#exit-address-family</code> | Exit VPNv6 address family mode. |
| <code>(config-router)#address-family ipv6 vrf IPI</code> | Enter the IPv6 address family for VRF IPI. |
| <code>(config-router-af)#neighbor 2001::2 remote-as 200</code> | Activate CE inside IPv6 address family for vrf IPI. |
| <code>(config-router-af)#neighbor 2001::2 activate</code> | Activate the 6VPE neighbor so that it can accept VPN IPv6 routes. |
| <code>(config-router-af)#redistribute connected</code> | Redistribute the connected route under address family IPv6 for VRF IPI. |
| <code>(config-router-af)#exit-address-family</code> | Exit IPv6 Address Family mode. |
| <code>(config-router)#exit</code> | Exit router mode. |
| <code>(config)#interface lo</code> | Enter interface mode. |
| <code>(config-if)#ip address 1.1.1.1/32 secondary</code> | Assign the IP address to loopback interface. |
| <code>(config-if)#exit</code> | Exit interface mode. |
| <code>(config)#router ldp</code> | Enter router LDP mode. |
| <code>(config-router)#router-id 1.1.1.1</code> | Set the router ID to IP address 1.1.1.1 |
| <code>(config-router)#transport-address ipv4 1.1.1.1</code> | Configure transport address as loopback address. |
| <code>(config-router)#targeted-peer ipv4 3.3.3.3</code> | Configure targeted peer. |
| <code>(config-router-targeted-peer)#exit</code> | Exit-targeted-peer-mode |
| <code>(config-router)#exit</code> | Exit router LDP mode. |
| <code>(config)#router rsvp</code> | Enter RSVP router mode. |
| <code>(config-router)#exit</code> | Exit router mode. |

| | |
|---|--|
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-rsvp | Enable RSVP in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#ip address 20.10.10.1/24 | Assign IP address to interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Enter OSPF router mode. |
| (config-router)#ospf router-id 1.1.1.1 | Configure OSPF router id same as loopback ip address. |
| (config-router)#network 1.1.1.1/32 area 0 (config-router)#network 20.10.10.1/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#rsvp-trunk toPE2 | Enter the trunk mode for RSVP. |
| (config-trunk)#to 3.3.3.3 | Specify IPv4 Egress for the LSP. |
| (config-trunk)#commit | Commit the transaction. |

P1

| | |
|--|--|
| #configure terminal | Enter configure mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 2.2.2.2/32 secondary | Assign the IP address to loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#router-id 2.2.2.2 | Set the router ID to IP address 2.2.2.2 |
| (config-router)#transport-address ipv4 2.2.2.2 | Configure transport address as loopback address. |
| (config-router)#exit | Exit router mode. |
| (config)#router rsvp | Enter RSVP router mode. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xe1 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-rsvp | Enable RSVP in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#ip address 20.10.10.2/24 | Assign IP address to interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface. |
| (config-if)#enable-rsvp | Enable RSVP in interface. |
| (config-if)#enable-ldp ipv4 | Enable ldp in interface. |

| | |
|---|--|
| (config-if)#ip address 20.10.20.1/24 | Assign IP address to interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Enter OSPF router mode. |
| (config-router)#ospf router-id 2.2.2.2 | Configure OSPF router id same as loopback ip address. |
| (config-router)#network 2.2.2.2/32 area 0 (config-router)#network 20.10.20.1/24 area 0 (config-router)#network 20.10.10.2/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#commit | Commit the transaction. |

PE2

| | |
|---|---|
| #configure terminal | Enter configure mode. |
| (config)#ip vrf IPI | Create a new VRF named IPI. |
| (config-vrf)#rd 1:101 | Assign the route distinguisher (RD) value as 1:101. |
| (config-vrf)#route-target both 100:200 | Import routes between route target (RT) ext-communities 100 and 200. |
| (config-vrf)#router-id 55.55.55.55 | Configure Router-id for VRF |
| (config-vrf)#exit | Exit VRF mode. |
| (config)#interface xe2 | Enter interface mode. |
| (config-if)#ip vrf forwarding IPI | Bind the interface connected to the CE router with VRF IPI. |
| (config-if)#ipv6 address 3001::1/64 | Assign the IPv6 address. |
| (config-if)#exit | Exit interface mode. |
| (config)#router bgp 100 | Enter BGP router mode. |
| (config-router)#bgp router-id 3.3.3.3 | Configure BGP router-id |
| (config-router)#neighbor 1.1.1.1 remote-as 100 | Configure 6VPE2 as an iBGP peer. |
| (config-router)#neighbor 1.1.1.1 update-source lo | Update the source as loopback for iBGP peering with the remote 6VPE router. |
| (config-router)#address-family ipv4 unicast | Enter address family mode |
| (config-router-af)#neighbor 1.1.1.1 activate | Activate the neighbor |
| (config-router-af)#exit-address-family | Exit address family mode |
| (config-router)#address-family vpnv6 unicast | Enter VPNv6 address family mode. |
| (config-router-af)#neighbor 1.1.1.1 activate | Activate the 6VPE neighbor so that it can accept VPN IPv6 routes. |
| (config-router-af)#exit-address-family | Exit VPNv6 address family mode. |
| (config-router)#address-family ipv6 vrf IPI | Enter the IPv6 address family for VRF IPI. |
| (config-router-af)#neighbor 3001::2 remote-as 300 | Activate CE inside IPv6 address family for vrf IPI. |

| | |
|--|---|
| (config-router-af)#neighbor 3001::2 activate | Activate the neighbor |
| (config-router-af)#redistribute connected | Redistribute the connected route under address family IPv6 for VRF IPI. |
| (config-router-af)#exit-address-family | Exit IPv6 Address Family mode. |
| (config-router)#exit | Exit router mode. |
| (config)#interface lo | Enter interface mode. |
| (config-if)#ip address 3.3.3.3/32 secondary | Assign the IP address to loopback interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ldp | Enter router LDP mode. |
| (config-router)#router-id 3.3.3.3 | Set the router ID to IP address 3.3.3.3 |
| (config-router)#transport-address ipv4 3.3.3.3 | Configure transport address as loopback address. |
| (config-router)#targeted-peer ipv4 1.1.1.1 | Configure targeted peer. |
| (config-router-targeted-peer)#exit | Exit-targeted-peer-mode |
| (config-router)#exit | Exit router mode |
| (config)#router rsvp | Enter RSVP router mode. |
| (config-router)#exit | Exit router mode. |
| (config)#interface xel | Enter interface mode. |
| (config-if)#label-switching | Enable label switching in interface |
| (config-if)#enable-rsvp | Enable RSVP in interface. |
| (config-if)#enable-ldp ipv4 | Enable LDP in interface. |
| (config-if)#ip address 20.10.20.2/24 | Assign IP address to interface. |
| (config-if)#exit | Exit interface mode. |
| (config)#router ospf 100 | Enter OSPF router mode. |
| (config-router)#ospf router-id 3.3.3.3 | Configure OSPF router id same as loopback ip address. |
| (config-router)#network 3.3.3.3/32 area 0 (config-router)#network 20.10.20.2/24 area 0 | Define the network on which OSPF runs and associate area id. |
| (config-router)#exit | Exit OSPF router mode. |
| (config)#rsvp-trunk toPE1 | Enter the trunk mode for RSVP. |
| (config-trunk)#to 1.1.1.1 | Specify IPv4 Egress for the LSP. |
| (config-trunk)#commit | Commit the transaction. |

Validation

CE1

```
CE1#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
```

```

E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,
N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C      ::1/128 via ::, lo, 01:38:28
C      2001::/64 via ::, xe1, 01:20:30
S      2ffe::/64 [1/0] via ::, xe1, 00:01:27
B      3001::/64 [20/0] via fe80::5054:ff:fe29:189d, xe1, 00:06:40
B      3ffe::/64 [20/0] via fe80::5054:ff:fe29:189d, xe1, 00:02:24
C      fe80::/64 via ::, xe3, 01:38:28

CE1#show ipv6 bgp summary vrf all
BGP router identifier 44.44.44.44, local AS number 200
BGP table version is 4
3 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv  MsgSen  TblVer  InQ   OutQ   Up/Down  State/PfxRcd
2001::1           4    100  1167    1522     4       0     0   00:13:23      3

Total number of neighbors 1

Total number of Established sessions 1

```

PE1

```

PE1#show ipv6 route vrf IPI
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "IPI"
C      2001::/64 via ::, xe1, 01:12:03
B      2ffe::/64 [20/0] via fe80::5054:ff:fe60:f4e5, xe1, 00:02:05
B      3001::/64 [200/0] via ::ffff:3.3.3.3, 00:08:02
B      3ffe::/64 [200/0] via ::ffff:3.3.3.3, 00:03:33
C      fe80::/64 via ::, xe1, 01:12:32

PE1#show ip bgp summary vrf all
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv  MsgSen  TblVer  InQ   OutQ   Up/Down  State/PfxRcd
1.1.1.1           4    100    78      84       1       0     0   00:32:15      0

Total number of neighbors 1

Total number of Established sessions 1

PE1#show ipv6 bgp summary vrf all
BGP router identifier 55.55.55.55, local AS number 100
BGP VRF IPI Route Distinguisher: 1:101
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv  MsgSen  TblVer  InQ   OutQ   Up/Down  State/PfxRcd

```



```
3001::2          4    200    79          81      1      0      0 00:32:15          2
```

Total number of neighbors 1

Total number of Established sessions 1

PE1#show mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
 B - BGP FTN, K - CLI FTN, t - tunnel, P - SR Policy FTN,
 L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
 U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

| Code | FEC | FTN-ID | Nhlfe-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|---------|----------|------------|------|----------|-----------|------|
| Intf | ELC | Nexthop | | | | | | |
| L> | 2.2.2.2/32 | | 1 | 1 | - | LSP_ | | |
| DEFAULT | 3 | xe2 | No | 20.10.10.2 | | | | |
| R(t)> | 3.3.3.3/32 | | 4 | 4 | 5001 | Yes | LSP_ | |
| DEFAULT | 24320 | xe2 | No | 20.10.10.2 | | | | |
| L | 3.3.3.3/32 | | 3 | 2 | - | LSP_ | | |
| DEFAULT | 24960 | xe2 | No | 20.10.10.2 | | | | |
| L> | 20.10.20.0/24 | | 2 | 1 | - | LSP_ | | |
| DEFAULT | 3 | xe2 | No | 20.10.10.2 | | | | |

PE1#show rsdp session

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
 State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
 * indicates the session is active with local repair at one or more nodes
 (P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

| To | From | Type | LSPName | State |
|---------|---------|---------|------------------|-----------------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 3.3.3.3 | 1.1.1.1 | PRI | toPE2-Primary | UP |
| | 24320 | DEFAULT | | 00:08:44 1 1 SE |

Total 1 displayed, Up 1, Down 0.

Egress RSVP:

| To | From | Type | LSPName | State |
|---------|---------|-------|------------------|--------------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 1.1.1.1 | 3.3.3.3 | PRI | toPE1-Primary | UP |
| SE | 24960 | - | ELSP_CON | 00:08:39 1 1 |

Total 1 displayed, Up 1, Down 0.

PE1#show mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM
 K - CLI ILM, T - MPLS-TP, s - Stitched ILM
 S - SNMP, L - LDP, R - RSVP, C - CRLDP
 B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
 O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
 P - SR Policy, U - unknown

| Code | FEC/VRF/L2CKT | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|----------|---------------|--------|----------|-----------|---------|-----------|
| Intf/VRF | Nexthop | | LSP-Type | | | |
| B> | IPI | 2 | 25600 | Nolabel | N/A | IPI |
| | LSP_DEFAULT | | | | | N/A |
| LT | | | | | | |
| R> | 3.3.3.3/32 | 2 | 24320 | Nolabel | N/A | N/A |
| | ELSP_CONF | | | | | 127.0.0.1 |
| IG | | | | | | |

PE1#show mpls vrf-table

Output for IPv6 VRF table with id: 2

Primary FTN entry with FEC: 3001::/64, id: 1, row status: Active, Tunnel-Policy: N/A
 Owner: BGP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
 Transport Tunnel id: 5001, Protected LSP id: 2201, QoS Resource id: 0, Description: N/A, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
 Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 5, owner: BGP, Stale: NO, BGP out intf: xe2, transport out intf: xe2, out label: 24320

Nexthop addr: 3.3.3.3 cross connect ix: 5, op code: Push and Lookup

Primary FTN entry with FEC: 3ffe::/64, id: 2, row status: Active, Tunnel-Policy: N/A
 Owner: BGP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
 Transport Tunnel id: 5001, Protected LSP id: 2201, QoS Resource id: 0, Description: N/A, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 5
 Owner: BGP, Persistent: NO, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 5, owner: BGP, Stale: NO, BGP out intf: xe2, transport out intf: xe2, out label: 24320

Nexthop addr: 3.3.3.3 cross connect ix: 5, op code: Push and Lookup

PE1#show mpls ftn-table

Primary FTN entry with FEC: 2.2.2.2/32, id: 1, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe2, out label: 3
 Nexthop addr: 20.10.10.2 cross connect ix: 1, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 4, row status: Active, Tunnel-Policy: N/A
 Owner: RSVP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 5001, Protected LSP id: 2201, QoS Resource id: 2, Description: toPE2, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 4
 Owner: RSVP, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 4, owner: RSVP, Stale: NO, out intf: xe2, out label: 24320
 Nexthop addr: 20.10.10.2 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 3.3.3.3/32, id: 3, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 2, in intf: - in label: 0 out-segment ix: 2
 Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Down
 Out-segment with ix: 2, owner: LDP, Stale: NO, out intf: xe2, out label: 24960
 Nexthop addr: 20.10.10.2 cross connect ix: 2, op code: Push

Primary FTN entry with FEC: 20.10.20.0/24, id: 2, row status: Active, Tunnel-Policy: N/A
 Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
 Tunnel id: 0, Protected LSP id: 0, Description: N/A, Color: 0
 Matched bytes:0, pkts:0, TX bytes:0, Pushed pkts:0
 Cross connect ix: 1, in intf: - in label: 0 out-segment ix: 1
 Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
 Out-segment with ix: 1, owner: N/A, Stale: NO, out intf: xe2, out label: 3
 Nexthop addr: 20.10.10.2 cross connect ix: 1, op code: Push

PE1#show ip bgp vpnv6 all

Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i - internal, l - labeled

S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

| Network | Next Hop | Metric | LocPrf | Weight | Path |
|--|----------------------------------|--------|--------|--------|------|
| Route Distinguisher: 1:100 (Default for VRF IPI) | | | | | |
| *> 1 2001::/64 | :: | 0 | 100 | 32768 | ? |
| * 2001::/64 | 2001::2(fe80::5054:ff:fe60:f4e5) | 0 | 100 | 0 200 | ? |
| *> 1 2ffe::/64 | 2001::2(fe80::5054:ff:fe60:f4e5) | 0 | 100 | 0 200 | ? |
| *>i 3001::/64 | ::ffff:3.3.3.3 | 0 | 100 | 0 | ? |

```

*>i 3ffe::/64      ::ffff:3.3.3.3      0      100      0      300 ?
Announced routes count = 3
Accepted routes count = 2
Route Distinguisher: 1:101
*>i 3001::/64      ::ffff:3.3.3.3      0      100      0      ?
*>i 3ffe::/64      ::ffff:3.3.3.3      0      100      0      300 ?
Announced routes count = 0
Accepted routes count = 2

PE1#show ip bgp neighbors
BGP neighbor is 3.3.3.3, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 1.1.1.1, remote router ID 3.3.3.3
  BGP state = Established, up for 00:09:55
  Last read 00:00:21, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
  Received 27 messages, 0 notifications, 0 in queue
  Sent 27 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

For address family: VPNv6 Unicast
  BGP table version 3, neighbor version 3
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  2 accepted prefixes
  2 announced prefixes

Connections established 1; dropped 0
Local host: 1.1.1.1, Local port: 33537
Foreign host: 3.3.3.3, Foreign port: 179
Nexthop: 1.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

BGP neighbor is 2001::2, vrf IPI, remote AS 200, local AS 100, external link
  BGP version 4, local router ID 77.77.77.77, remote router ID 44.44.44.44
  BGP state = Established, up for 00:16:19
  Last read 00:00:10, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv6 Unicast: advertised and received
  Received 42 messages, 0 notifications, 0 in queue
  Sent 42 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
For address family: IPv6 Unicast
  BGP table version 1, neighbor version 1
  Index 0, Offset 0, Mask 0x1
  Community attribute sent to this neighbor (standard)
  2 accepted prefixes
  3 announced prefixes

Connections established 1; dropped 0
Local host: 2001::1, Local port: 34776
Foreign host: 2001::2, Foreign port: 179
Nexthop: 77.77.77.77
Nexthop global: 2001::1

```

```
Nexthop local: fe80::5054:ff:fe29:189d
BGP connection: shared network
```

P1

```
P1#show mpls forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN,
        B - BGP FTN, K - CLI FTN, t - tunnel
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
```

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|------------|--------|------------|-----|----------|-----------|------|
| Intf | Nexthop | | | | | | |
| L> | 1.1.1.1/32 | 1 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe2 | 20.10.10.1 | | | | |
| L> | 3.3.3.3/32 | 2 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe1 | 20.10.20.2 | | | | |

```
P1#show mpls ilm-table
```

```
Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM, T - MPLS-TP, S - Stitched ILM
        S - SNMP, L - LDP, R - RSVP, C - CRLDP
        B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
        O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
        U - unknown
```

| Code | FEC/VRF | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|------|-------------|--------|----------|-----------|---------|----------------|
| Intf | Nexthop | | LSP-Type | | | |
| R> | 1.1.1.1/32 | 2 | 24321 | 24960 | N/A | xe1 20.10.10.1 |
| | ELSP_CONFIG | | | | | |
| R> | 3.3.3.3/32 | 1 | 24320 | 24960 | N/A | xe2 20.10.20.2 |
| | ELSP_CONFIG | | | | | |
| L> | 1.1.1.1/32 | 4 | 24961 | 3 | N/A | xe1 20.10.10.1 |
| | LSP_DEFAULT | | | | | |
| L> | 3.3.3.3/32 | 5 | 24960 | 3 | N/A | xe2 20.10.20.2 |
| | LSP_DEFAULT | | | | | |

```
P1#show ip ospf neighbor
```

```
Total number of full neighbors: 2
```

```
OSPF process 100 VRF(default):
```

| Neighbor ID | Pri | State | Dead Time | Address | Interface | Instance ID |
|-------------|-----|-------------|-----------|------------|-----------|-------------|
| 1.1.1.1 | 1 | Full/Backup | 00:00:31 | 20.10.10.1 | xe1 | 0 |
| 3.3.3.3 | 1 | Full/DR | 00:00:32 | 20.10.20.2 | xe2 | 0 |

PE2

```
PE2#show ipv6 route vrf IPI
```

```
IPv6 Routing Table
```

```
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
        IA - OSPF inter area, E1 - OSPF external type 1,
        E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,
        N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
```

```
Timers: Uptime
```

```
IP Route Table for VRF "IPI"
```

```
C    ::1/128 via ::, lo.IPI, 00:24:23
C    3001::/64 via ::, xe1, 00:24:22
B    3ffe::/64 [20/0] via fe80::3617:ebff:fe0e:1201, xe1, 00:05:28
C    fe80::/64 via ::, xe1, 00:24:22
```

```
PE2#show ip bgp summary vrf all
```

```
BGP router identifier 55.55.55.55, local AS number 100
```

```

BGP VRF IPI Route Distinguisher: 1:100
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|----------|---|-----|--------|--------|--------|-----|------|----------|--------------|
| 3001::2 | 4 | 300 | 116 | 181 | 1 | 0 | 0 | 00:22:05 | 0 |

Total number of neighbors 1

```

Total number of Established sessions 1
BGP router identifier 3.3.3.3, local AS number 100
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

```

| Neighbor | V | AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd |
|----------|---|-----|--------|--------|--------|-----|------|----------|--------------|
| 1.1.1.1 | 4 | 100 | 65 | 66 | 1 | 0 | 0 | 00:26:21 | 0 |

Total number of neighbors 1

Total number of Established sessions 1

```
PE2#show mpls forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN,
        B - BGP FTN, K - CLI FTN, t - tunnel
        L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
        U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

```

| Code | FEC | FTN-ID | Tunnel-id | Pri | LSP-Type | Out-Label | Out- |
|---------|---------------|--------|------------|-----|----------|-----------|------|
| Intf | Nexthop | | | | | | |
| R(t)> | 1.1.1.1/32 | 1 | 5001 | Yes | LSP_ | | |
| DEFAULT | 24321 | xe2 | 20.10.20.1 | | | | |
| L | 1.1.1.1/32 | 2 | 0 | Yes | LSP_ | | |
| DEFAULT | 24961 | xe2 | 20.10.20.1 | | | | |
| L> | 2.2.2.2/32 | 3 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe2 | 20.10.20.1 | | | | |
| L> | 20.10.10.0/24 | 4 | 0 | Yes | LSP_ | | |
| DEFAULT | 3 | xe2 | 20.10.20.1 | | | | |

```
PE2#show rsvp session
```

```

Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes

```

```
Ingress RSVP:
```

| To | From | Type | LSPName | State |
|---------|---------|---------|------------------|-----------------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 1.1.1.1 | 3.3.3.3 | PRI | toPE1-Primary | UP |
| | 24321 | DEFAULT | | 00:23:21 1 1 SE |

Total 1 displayed, Up 1, Down 0.

```
Egress RSVP:
```

| To | From | Type | LSPName | State |
|---------|---------|-------|------------------|--------------|
| Uptime | Rt | Style | Labelin Labelout | DSType |
| 3.3.3.3 | 1.1.1.1 | PRI | toPE2-Primary | UP |
| SE | 24960 | - | ELSP_CON | 00:23:33 1 1 |

Total 1 displayed, Up 1, Down 0.

```
PE2#show mpls ilm-table
```

```

Codes: > - installed ILM, * - selected ILM, p - stale ILM
        K - CLI ILM, T - MPLS-TP, S - Stitched ILM
        S - SNMP, L - LDP, R - RSVP, C - CRLDP
        B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
        O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
        U - unknown

```

| Code | FEC/VRF | ILM-ID | In-Label | Out-Label | In-Intf | Out- |
|------|---------|--------|----------|-----------|---------|------|
| Intf | Nexthop | | LSP-Type | | | |

```

B> 3ffe::/64          3          24321      N/A        N/A        xe1        fe80::3617:ebff
:fe0e:1201 LSP_DEFAULT
B> 3001::/64          2          24320      N/A        N/A        xe1        ::
LSP_DEFAULT
R> 3.3.3.3/32         1          24960      N/A        N/A        N/A        127.0.0.1
ELSP_CONFIG

```

```
PE2#show ip bgp neighbors
```

```

BGP neighbor is 1.1.1.1, remote AS 100, local AS 100, internal link
  BGP version 4, local router ID 3.3.3.3, remote router ID 1.1.1.1
  BGP state = Established, up for 00:23:39
  Last read 00:00:27, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family VPNv6 Unicast: advertised and received
  Received 58 messages, 0 notifications, 0 in queue
  Sent 60 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
  Update source is lo
For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

```

```

For address family: VPNv6 Unicast
  BGP table version 4, neighbor version 4
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)
  2 accepted prefixes
  2 announced prefixes

```

```

Connections established 1; dropped 0
Local host: 3.3.3.3, Local port: 37145
Foreign host: 1.1.1.1, Foreign port: 179
Nexthop: 3.3.3.3
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network

```

```

BGP neighbor is 3001::2, vrf IPI, remote AS 300, local AS 100, external link
  BGP version 4, local router ID 55.55.55.55, remote router ID 66.66.66.66
  BGP state = Established, up for 00:19:23
  Last read 00:00:05, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 110 messages, 0 notifications, 0 in queue
  Sent 113 messages, 62 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 30 seconds
For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  0 accepted prefixes
  0 announced prefixes

For address family: IPv6 Unicast
  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (standard)
  2 accepted prefixes
  3 announced prefixes

```

```

Connections established 1; dropped 0
Local host: 3001::1, Local port: 179
Foreign host: 3001::2, Foreign port: 58741
NextHop: 55.55.55.55
NextHop global: 3001::1
NextHop local: fe80::da9e:f3ff:fec9:65a1
BGP connection: shared network
Last Reset: 00:19:28, due to OPEN Message Error (Notification sent)
Notification Error Message: (OPEN Message Error/Bad BGP Identifier.)

PE2#show ip bgp vpnv6 all
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric      LocPrf      Weight Path
Route Distinguisher: 1:100 (Default for VRF IPI)
*>i  2001::/64      ::ffff:101:101          0           100          0          ?
*>i  2ffe::/64       ::ffff:101:101          0           100          0          200 ?
*> 1 3001::/64      ::                0           100          32768      ?
*   3001::/64       3001::2(fe80::3617:ebff:fe0e:1201)
                                   0           100          0           300 ?
*> 1 3ffe::/64      3001::2(fe80::3617:ebff:fe0e:1201)
                                   0           100          0           300 ?

Announced routes count = 3
Accepted routes count = 2
Route Distinguisher: 1:100
*>i  2001::/64      ::ffff:101:101          0           100          0          ?
*>i  2ffe::/64       ::ffff:101:101          0           100          0          200 ?
Announced routes count = 0
Accepted routes count = 2

PE2#show mpls ftn-table
Primary FTN entry with FEC: 1.1.1.1/32, id: 1, row status: Active
Owner: RSVP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 5001, Protected LSP id: 2201, QoS Resource id: 2, Description: toPE1
Cross connect ix: 4, in intf: - in label: 0 out-segment ix: 3
Owner: RSVP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 3, owner: RSVP, out intf: xel, out label: 24321
NextHop addr: 20.10.20.1 cross connect ix: 4, op code: Push

Primary FTN entry with FEC: 1.1.1.1/32, id: 2, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
Cross connect ix: 5, in intf: - in label: 0 out-segment ix: 4
Owner: LDP, Persistent: No, Admin Status: Down, Oper Status: Down
Out-segment with ix: 4, owner: LDP, out intf: xel, out label: 24961
NextHop addr: 20.10.20.1 cross connect ix: 5, op code: Push

Primary FTN entry with FEC: 2.2.2.2/32, id: 3, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 5
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 5, owner: LDP, out intf: xel, out label: 3
NextHop addr: 20.10.20.1 cross connect ix: 6, op code: Push

Primary FTN entry with FEC: 20.10.10.0/24, id: 4, row status: Active
Owner: LDP, distance: 0, Action-type: Redirect to LSP, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 0, Protected LSP id: 0, Description: N/A
Cross connect ix: 6, in intf: - in label: 0 out-segment ix: 5
Owner: LDP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 5, owner: LDP, out intf: xel, out label: 3
NextHop addr: 20.10.20.1 cross connect ix: 6, op code: Push

```

```

PE2#show mpls vrf-table
Output for IPv6 VRF table with id: 2
Primary FTN entry with FEC: 2001::/64, id: 1, row status: Active
Owner: BGP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 5001, Protected LSP id: 2201, QoS Resource id: 0, Description: N/A
Cross connect ix: 3, in intf: - in label: 0 out-segment ix: 2
Owner: BGP, Persistent: NO, Admin Status: Up, Oper Status: Up
Out-segment with ix: 2, owner: BGP, out intf: N/A, out label: 24320
Nexthop addr: 1.1.1.1 cross connect ix: 3, op code: Push and Lookup

Primary FTN entry with FEC: 2ffe::/64, id: 2, row status: Active
Owner: BGP, distance: 0, Action-type: Redirect to Tunnel, Exp-bits: 0x0, Incoming DSCP: none
Tunnel id: 5001, Protected LSP id: 2201, QoS Resource id: 0, Description: N/A
Cross connect ix: 7, in intf: - in label: 0 out-segment ix: 6
Owner: BGP, Persistent: No, Admin Status: Up, Oper Status: Up
Out-segment with ix: 6, owner: BGP, out intf: N/A, out label: 24321
Nexthop addr: 1.1.1.1 cross connect ix: 7, op code: Push and Lookup

```

CE2

```

CE2#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, I - IS-IS, B - BGP
Timers: Uptime

IP Route Table for VRF "default"
C    ::1/128 via ::, lo, 00:37:26
B    2001::/64 [20/0] via fe80::da9e:f3ff:fec9:65a1, xe2, 00:20:44
B    2ffe::/64 [20/0] via fe80::da9e:f3ff:fec9:65a1, xe2, 00:09:52
C    3001::/64 via ::, xe2, 00:27:07
S    3ffe::/64 [1/0] via ::, xe2, 00:07:31
C    fe80::/64 via ::, xe2, 00:37:26

CE2#show ip bgp summary vrf all
BGP router identifier 66.66.66.66, local AS number 300
BGP table version is 1
3 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS  MsgRcv  MsgSen  TblVer   InQ   OutQ   Up/Down  State/PfxRcd
3001::1           4    100    178    176      1      0      0  00:20:51      0

Total number of neighbors 1

Total number of Established sessions 1

```


Bridge Virtual Interface (BVI) Over 6vPE

Overview

BVI over 6vPE enables Layer 3 gateway functionality within an IPv6 L3VPN (6vPE) environment by bridging IPv6-capable Layer 2 domains to an MPLS-enabled IPv6 core. A Bridge Virtual Interface (BVI) serves as the logical routed interface that connects a local bridge domain (Layer 2) to the IPv6 VPN routing instance.

For details on BVI Over L3VPN, see [Bridge Virtual Interface \(BVI\) Over L3VPN \(page 2540\)](#). For QoS and ACL configuration for BVIs, see the *QoS and ACL over Bridge Virtual Interface (BVI)* chapter in Layer 3 guide.

Prerequisites

- Define Interfaces and Loopback Addresses:

Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish efficient network routing and communication.

```
interface lo
  ip address 127.0.0.1/8
  ip address 135.1.1.27/32 secondary
  ipv6 address ::1/128

interface po6
  ip address 10.1.1.1/30
interface xe6
  channel-group 6 mode active
```

- Configure IGP for Dynamic Routing:

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

- ISIS Configuration:

```
router isis 1
  is-type level-2-only
  metric-style wide
  mpls traffic-eng router-id 10.12.183.1
  mpls traffic-eng level-2
  capability cspf
  dynamic-hostname
  fast-reroute ti-lfa level-2 proto ipv4
  net 49.0000.0000.0027.00
  passive-interface lo

interface po6
  isis network point-to-point
  ip router isis 1
```

- OSPF Configuration:

```
router ospf 100
  ospf router-id 10.12.183.1
  network 10.12.183.1/32 area 0.0.0.0
  network 10.1.1.0/24 area 0.0.0.0
  network 10.1.1.0/24 area 0.0.0.0
```

- Configure LDP for Label Transport:
Configure label-switching on all the nodes to help exchange the network packets at a lower lever rather than the traditional network layer and enable ldp to transport those labels.

```
router ldp
targeted-peer ipv4 10.12.183.3
exit-targeted-peer
transport address ipb4 10.12.183.1

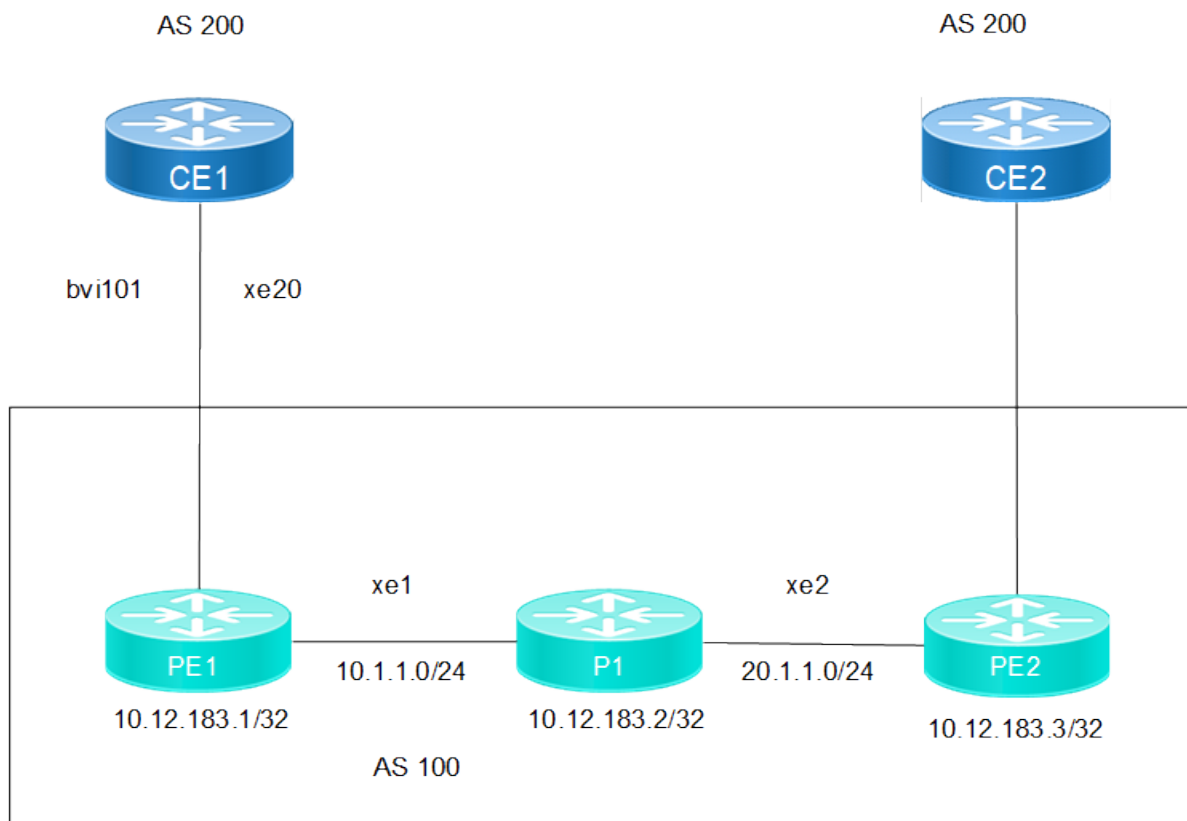
Interface po6
label-switching
enable-ldp ipv4
```

Configuration

Topology

The sample topology includes Edge nodes (PE1 and PE2), core Node (P1), and Customer Edge nodes (CE1 and CE2).

Figure 106. BVI Interface Topology



Note: Before configuring BVI-6vPE, meet all the pre-requisite for the following nodes:

- Edge nodes: PE1 and PE2
- Core nodes: P1, P3, and P4

Configure CE1 Router

1. Set up the interface with VLAN encapsulation and assign an IP address.

```
CE1(config)#interface xe20.101
CE1(config-if)#encapsulation dot1q 101
CE1(config-if)#ipv6 address 2001::2/64
CE1(config-if)#commit
CE1(config-if)#exit
```

2. BGP Configuration with a neighbor in the same AS (200).

```
CE1(config)# router bgp 200
CE1(config-router)# neighbor 2001::1 remote-as 100
CE1(config-router)#address-family ipv4 unicast
CE1(config-router-af)#neighbor 2001::1 activate
CE1(config-router-af)#exit-address-family
CE1(config-router)#commit
CE1(config-router)#exit
```

Configure PE1 Router

1. Create a VRF instance (VRF101) and define route targets.

```
PE1(config)# ip vrf VRF101
PE1(config-vrf)# rd 10.12.183.1:100
PE1(config-vrf)# route-target both 100:101
PE1(config-vrf)# commit
PE1(config-vrf)# exit
```

2. Configure the access port with VLAN encapsulation and enable rewriting.

```
PE1(config)#interface xe20.101 switchport
PE1(config-if)#encapsulation dot1q 101
PE1(config-if)#rewrite pop
PE1(config-if)# commit
PE1(config-if)# exit
```

3. Set up a BVI for VRF forwarding.

```
PE1(config)# interface bvi101
PE1(config-if)#ip vrf forwarding VRF101
PE1(config-if)#ipv6 address 2001::1/64
PE1(config-if)# commit
PE1(config-if)# exit
```

4. Configure BGP with internal (PE) and external (CE) neighbors.

```
PE1(config)# router bgp 100
PE1(config-router)# neighbor 10.12.183.3 remote-as 100
PE1(config-router)# neighbor 10.12.183.3 update-source lo
PE1(config-router)#address-family ipv4 unicast
PE1(config-router-af)#neighbor 10.12.183.3 activate
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family vpnv4 unicast
PE1(config-router-af)#neighbor 10.12.183.3 activate
PE1(config-router-af)#exit-address-family
```

5. Enable VPNv4 and redistribute connected routes.

```
PE1(config-router)#address-family ipv4 vrf VRF101
PE1(config-router-af)#redistribute connected
PE1(config-router-af)#neighbor 2001::2 remote-as 200
PE1(config-router-af)#neighbor 2001::2 activate
PE1(config-router-af)#exit-address-family
PE1(config-router)#commit
PE1(config-router)#exit
```

6. Set up a bridge domain to associate the interface with BVI.

```
PE1(config)# bridge-domain 101
PE1(config-bridge-domain)#interface xe20.101
PE1(config-bridge-domain)#routed-interface bv101
PE1(config-bridge-domain)#commit
PE1(config-bridge-domain)#exit
```

Configure PE2 Router

1. Create the same VRF instance (VRF101) with route targets.

```
PE2(config)# ip vrf VRF101
PE2(config-vrf)# rd 10.12.183.3:100
PE2(config-vrf)# route-target both 100:101
PE2(config-vrf)#commit
PE2(config-vrf)#exit
```

2. Configure the access port with VLAN encapsulation and VRF forwarding.

```
PE2(config)# interface xe30.101
PE2(config-if)#encapsulation dot1q 101
PE2(config-if)#ip vrf forwarding VRF101
PE2(config-if)#ipv6 address 3001::1/64
PE2(config-if)# commit
PE2(config-if)# exit
```

3. Assign an IP address to the interface.

```
PE2(config)# router bgp 100
PE2(config-router)# neighbor 10.12.183.1 remote-as 100
PE2(config-router)# neighbor 10.12.183.1 update-source lo
PE2(config-router)#address-family ipv4 unicast
PE2(config-router-af)#neighbor 10.12.183.1 activate
PE2(config-router-af)#exit-address-family
```

4. Configure BGP with internal (PE) and external (CE) neighbors.

```
PE2(config-router)#address-family vpnv4 unicast
PE2(config-router-af)#neighbor 10.12.183.1 activate
PE2(config-router-af)#exit-address-family
```

5. Enable VPNv4 and redistribute connected routes.

```
PE2(config-router)#address-family ipv4 vrf VRF101
PE2(config-router-af)#redistribute connected
PE2(config-router-af)#neighbor 3001::2 remote-as 200
PE2(config-router-af)#neighbor 3001::2 activate
PE2(config-router-af)#exit-address-family
PE2(config-router)#commit
PE2(config-router)#exit
```

Configure CE2 Router

1. Set up the interface with VLAN encapsulation and assign an IP address.

```
CE2#configure terminal
CE2(config)#interface xe30.101
CE2(config-if)#encapsulation dot1q 101
CE2(config-if)#ipv6 address 3001::2/64
CE2(config-if)#commit
CE2(config-if)#exit
```

2. Configure BGP with a neighbor in the same AS (200).

```
CE2(config)# router bgp 200
CE2(config-router)#neighbor 3001::1 remote-as 100
CE2(config-router)#address-family ipv4 unicast
CE2(config-router-af)#neighbor 3001::1 activate
```

```
CE2(config-router-af)#exit-address-family
CE2(config-router)#commit
CE2(config-router)#exit
```

Running Configuration

CE1

```
!
interface xe20.101
 encapsulation dot1q 101
 ipv6 address 2001::2/64
 exit
!
router bgp 200
 neighbor 2001::1 remote-as 100
 address-family ipv4 unicast
  neighbor 2001::1 activate
 exit-address-family
 exit
!
```

PE1

PE1(config)# ip vrf VRF101

```
!
ip vrf VRF101
 rd 10.12.183.1:100
 route-target both 100:101
 exit
!
interface xe20.101 switchport
 encapsulation dot1q 101
 rewrite pop
 exit
!
interface bvi101
 ip vrf forwarding VRF101
 ipv6 address 2001::1/64
 exit
!
router bgp 100
 neighbor 10.12.183.3 remote-as 100
 neighbor 10.12.183.3 update-source lo
 address-family ipv4 unicast
  neighbor 10.12.183.3 activate
 exit-address-family
 address-family vpnv4 unicast
  neighbor 10.12.183.3 activate
 exit-address-family
 address-family ipv4 vrf VRF101
 redistribute connected
 neighbor 2001::2 remote-as 200
 neighbor 2001::2 activate
 exit-address-family
 exit
!
bridge-domain 101
 interface xe20.101
 router-interface bvi101
 exit
!
```

PE2

```
!  
ip vrf VRF101  
  rd 10.12.183.3:100  
  route-target both 100:101  
  exit  
!  
interface xe30.101  
  encapsulation dot1q 101  
  ip vrf forwarding VRF101  
  ipv6 address 3001::1/64  
  exit  
!  
router bgp 100  
  neighbor 10.12.183.1 remote-as 100  
  neighbor 10.12.183.1 update-source lo  
  address-family ipv4 unicast  
    neighbor 10.12.183.1 activate  
  exit-address-family  
  address-family vpnv4 unicast  
    neighbor 10.12.183.1 activate  
  exit-address-family  
  address-family ipv4 vrf VRF101  
    redistribute connected  
    neighbor 3001::2 remote-as 200  
    neighbor 3001::2 activate  
  exit-address-family  
  exit  
!
```

CE2

```
#show running-config  
!  
feature netconf-ssh vrf management  
feature netconf-tls vrf management  
no feature netconf-ssh  
no feature netconf-tls  
!  
service password-encryption  
!  
logging console 5  
snmp-server enable traps link linkDown  
snmp-server enable traps link linkUp  
!  
hardware-profile statistics ingress-acl enable  
hardware-profile statistics ac-lif enable  
!  
qos enable  
!  
hostname PE2  
no ip domain-lookup  
ip domain-lookup vrf management  
ip name-server vrf management 10.12.3.23  
bridge 1 protocol ieee vlan-bridge  
tfo Disable  
errdisable cause stp-bpdu-guard  
no feature telnet vrf management  
no feature telnet  
feature ssh vrf management  
no feature ssh  
snmp-server enable snmp vrf management  
snmp-server view all .1 included vrf management  
snmp-server community test vrf management  
feature dns relay  
ip dns relay
```

```
ipv6 dns relay
feature ntp vrf management
ntp enable vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 10
!
router-id 10.12.183.3
!
ip vrf management
!
ip vrf VRF101
  rd 10.12.183.3:100
  route-target both 100:101
!
router ldp
  fast-reroute
  session-protection duration 40
  targeted-peer ipv4 10.12.183.1
  exit-targeted-peer-mode
  transport-address ipv4 10.12.183.3
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface lo
  ip address 127.0.0.1/8
  ip address 10.12.183.3/32 secondary
  ipv6 address ::1/128
  ip router isis ISIS-IGP-100
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
interface xe2
  ip address 20.1.1.2/24
  label-switching
  ip router isis ISIS-IGP-100
  enable-ldp ipv4
!
interface xe30
!
interface xe30.101
  description L3VPN-VRF101
  encapsulation dot1q 101
  ip vrf forwarding VRF101
  ip address 103.0.1.1/24
!
exit
!
router ospf 100
  ospf router-id 10.12.183.3
  network 10.12.183.3/32 area 0.0.0.0
  network 20.1.1.0/24 area 0.0.0.0
!
router isis ISIS-IGP-100
  is-type level-1
  metric-style wide
  mpls traffic-eng router-id 10.12.183.3
  mpls traffic-eng level-1
  capability cspf
```

```

dynamic-hostname
fast-reroute per-prefix remote-lfa level-1 proto ipv4 tunnel mpls-ldp
bfd all-interfaces
net 49.0001.0000.0000.0003.00
passive-interface lo
!
router bgp 100
neighbor 10.12.183.1 remote-as 100
neighbor 10.12.183.1 update-source lo
!
address-family ipv4 unicast
neighbor 10.12.183.1 activate
exit-address-family
!
address-family vpnv4 unicast
neighbor 10.12.183.1 activate
exit-address-family
!
address-family ipv4 vrf VRF101
redistribute connected
neighbor 103.0.1.2 remote-as 200
neighbor 103.0.1.2 activate
exit-address-family
!
exit
!
end

```

CE2

```

!
interface xe30.101
encapsulation dot1q 101
ipv6 address 3001::2/64
exit
!
router bgp 200
neighbor 3001::1 remote-as 100
address-family ipv4 unicast
neighbor 3001::1 activate
exit-address-family
exit
!

```

Validation

PE1

To Verify the L2 interfaces and bvi interface attached on a bridge domain

```

PE1#show running-config bridge-domain

bridge-domain 101
Interface xe20.101
routed-interface bvi101

```

To Verify the link status of L2 Interfaces attached to bridge-domain:

```

PE1#show bridge-domain

Total number of bridge-domains Configured: 1

Bridge Id  interfaces      Status
-----+-----+-----

```



```
101          xe20.101          UP
```

To Verify the link status of the BVI interface:

```
PE1#show ip interface bvi101 brief
```

```
'*' - address is assigned by dhcp client
```

| Interface | IP-Address | Admin-Status | Link-Status |
|-----------|------------|--------------|-------------|
| bvi101 | unassigned | up | up |

To Verify BGP session between PE1-CE1

```
PE1#show ipv6 bgp summary
```

```
BGP router identifier 10.12.183.1, local AS number 100
```

```
BGP table version is 11
```

```
1 BGP AS-PATH entries
```

```
0 BGP community entries
```

| Neighbor | V AS | MsgRcv | MsgSen | TblVer | InQ | OutQ | Up/Down | State/PfxRcd | Desc |
|----------|------|--------|--------|--------|-----|------|------------|--------------|------|
| 2001::2 | 4 | 100 | 280 | 283 | 11 | 0 | 0 00:18:47 | 0 | |

```
Total number of neighbors 1
```

```
Total number of Established sessions 1
```

VRF Forwarding table

```
PE1-7017#show mpls vrf-forwarding-table
```

```
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
```

```
(m) - Service mapped over multipath transport
```

```
(e) - Service mapped over ECMP
```

```
(D) - Down
```

```
Ext-Color - Extended-community color advertised by BGP
```

```
B(x) - BGP EVPN MPLS Services
```

| Code | FEC | FTN-ID | VRF-ID | Nhlfe-ID | Pri | Out-Label | Out- |
|------|-------------|--------|-----------|----------|-----|-----------|------|
| Intf | Nexthop | UpTime | Ext-Color | | | | |
| B> | 3001::/64 | 2 | 2 | 9 | - | - | - |
| | 00:14:55 | - | | 8 | Yes | 26240 | - |
| | 10.12.183.3 | - | - | | | | |

To Verify vrf Ping between PE1-CE1

```
PE1-7017#ping mpls l3vpn VRF101 3001::/64 detail
```

```
Sending 5 MPLS Echos to 3001::, timeout is 5 seconds
```

```
Codes:
```

```
'!' - Success, 'Q' - request not sent, '.' - timeout,
```

```
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
```

```
'N' - LBL Mapping Err, 'D' - DS Mismatch,
```

```
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
```

```
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
```

```
'P' - Protocol Error, 'X' - Unknown code,
```

```
'Z' - Reverse FEC Validation Failed
```

```
Type 'Ctrl+C' to abort
```

```
! seq_num = 1 10.1.1.2 1.11 ms
```

```
! seq_num = 2 10.1.1.2 0.62 ms
```

```
! seq_num = 3 10.1.1.2 0.51 ms
```

```
! seq_num = 4 10.1.1.2 0.59 ms
```

```
! seq_num = 5 10.1.1.2 0.76 ms
```

```
Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.51/0.81/1.11
```

To Verify the Ping between PE1-PE2

```
PE1-7017#ping mpls l3vpn VRF101 3001::/64 detail
Sending 5 MPLS Echos to 3001::, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

Type 'Ctrl+C' to abort

! seq_num = 1 10.1.1.2 1.11 ms
! seq_num = 2 10.1.1.2 0.62 ms
! seq_num = 3 10.1.1.2 0.51 ms
! seq_num = 4 10.1.1.2 0.59 ms
! seq_num = 5 10.1.1.2 0.76 ms

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.51/0.81/1.11
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

| Key Terms/Acronym | Description |
|-------------------|--|
| BUM | Broadcast, Unknown, Multicast |
| BVI | Bridge Virtual Interface |
| L3VPN | Layer 3 Virtual Private Network |
| IGP | Interior Gateway Protocol |
| ISIS | Intermediate System to Intermediate System |
| OSPF | Open Shortest Path First |
| LDP | Label Distribution Protocol |

APPENDIX A: SERVICE TEMPLATE CONFIGURATION

Overview

The Service Template mechanism in OcNOS (supported up to version 6.6.0) was introduced to simplify Layer 2 VPN (L2VPN) service deployment by decoupling traffic classification logic (such as VLAN matching) from physical interface configurations. A service-template could be applied to multiple interfaces, enabling consistent and reusable provisioning across access ports in VPWS, VPLS, and other L2VPN scenarios.



Note: Starting with OcNOS version 7.0.0, sub-interfaces replace Service Templates as the preferred configuration method. While legacy Service Template configurations remain functional, the associated commands are now hidden and will be completely removed in a future release. To ensure future-proofing and continued manageability, please migrate all existing Service Templates to sub-interfaces.

Purpose and Function

Service Templates enabled:

- Centralized VLAN matching and classification logic that could be applied across multiple interfaces.
- A method to bind L2VPN services (e.g., VPWS or VPLS) to access interfaces based on tag criteria, such as untagged, single-tagged, or double-tagged (Q-in-Q) traffic.
- Integration with:
 - `mpls-l2-circuit` for VPWS
 - `mpls-vpls` or `vpls-instance` for VPLS
 - Under the context of `switchport` interfaces

Typical Use Cases

Service Templates were commonly used in the following scenarios:

- **Virtual Private Wire Service (VPWS):** For point-to-point pseudowire configurations between two PE routers.
- **Virtual Private LAN Service (VPLS):** For multipoint Layer 2 VPN services spanning multiple sites.
- **VPLS Service Mapping:** Including static, LDP-signaled, and BGP-signaled VPLS instances using various tag matching schemes.

Common Syntax Patterns

This section provides typical examples of Service Template-based configurations used in VPLS and VPWS deployments in OcNOS versions up to 6.6.0. These patterns demonstrate how various service scenarios were modeled using service templates, including handling of untagged traffic, VLAN-tagged traffic, and VLAN tag rewrites.

VPLS – Untagged Interface Configuration

In this example, a basic service template without any match condition is used to identify untagged traffic on the access port.

```
service-template <service template name>
!
mpls vpls <vpls instance name> <vpls id>
  signaling ldp
  vpls-type <encapsulation type>
  vpls-peer <far end peer address>
  exit-signaling
exit-vpls
!
interface xe21
  switchport
  mpls-vpls <vpls instance name> service-template <service template name>
  exit-if-vpls
!
```

VPLS – VLAN Tagged Interface Configuration

Here, the service template matches traffic based on a single VLAN tag (outer VLAN). This is used for identifying tagged access traffic.

```
service-template <service template name>
  match outer-vlan <vlan tag>
!
mpls vpls <vpls instance name> <vpls id>
  signaling ldp
  vpls-type <encapsulation type>
  vpls-peer <far end peer address>
  exit-signaling
exit-vpls
!
interface xe21
  switchport
  mpls-vpls <vpls instance name> service-template <service template name>
  exit-if-vpls
!
```

VPLS – VLAN Tagged Interface with Rewrite

This example includes both a VLAN tag match and a rewrite directive. It translates the ingress VLAN tag to a new value before forwarding.

```

service-template <service template name>
  match outer-vlan <vlan tag>
  rewrite ingress translate <new vlan tag value to change to> outgoing-tpid dot1.q
!
mpls vpls <vpls instance name> <vpls id>
  signaling ldp
  vpls-type <encapsulation type>
  vpls-peer <far end peer address>
  exit-signaling
  exit-vpls
!
interface xe21
  switchport
  mpls-vpls <vpls instance name> service-template <service template name>
  exit-if-vpls
!

```

VPWS – Untagged Interface Configuration

This is a standard point-to-point L2 circuit configuration for untagged traffic using a service template.

```

service-template <service template name>
  match untagged
!
mpls l2-circuit <vpws name> <virtual circuit ID number> <far end neighbor>
!
interface xe21
  switchport
  mpls-l2-circuit <vpws name> service-template <service template name> primary
!

```

VPWS – VLAN Tagged Interface Configuration

This example uses a VLAN match to associate tagged ingress traffic with a VPWS instance.

```

service-template <service template name>
  match outer-vlan <vlan tag>
!
mpls l2-circuit <vpws name> <virtual circuit ID number> <far end neighbor>
!
interface xe21
  switchport
  mpls-l2-circuit <vpws name> service-template <service template name> primary
!

```

VPWS – VLAN Tagged Interface with Rewrite

This pattern includes VLAN tag rewriting on ingress, which is useful in service provider environments requiring normalization or mapping of customer VLANs.

```

service-template <service template name>
  match outer-vlan <vlan tag>
  rewrite ingress translate <new vlan tag value to change to> outgoing-tpid dot1.q
!
mpls l2-circuit <vpws name> <virtual circuit ID number> <far end neighbor>
!
interface xe21

```

```
switchport
mpls-l2-circuit <vpws name> service-template <service template name> primary
!
```

Documentation Reference

For detailed service-template configuration procedures in different service types, refer to the following sections in the [Layer 2 Virtual Private Network Configuration](#) in the OcNOS 6.5.x MPLS document:

- [Virtual Private LAN Service Configuration](#)
- [Static VPLS Configuration](#)
- [BGP-VPLS Configuration](#)
- [Static VPLS Service Mapping Configuration](#)
- [LDP-VPLS Service Mapping Configuration](#)
- [BGP-VPLS Service Mapping Configuration](#)
- [Virtual Private Wire Service Configuration](#)

Migration from Service Template to Sub-interface Configuration

Applicability: OcNOS Release 6.6.1 and later

Starting with OcNOS version 6.6.1, we would like to recommend using the sub-interface-based configuration model for MPLS Layer 2 VPN services, including VPLS and VPWS. We are phasing out the support for the traditional Service Template based configurations and would like to provide a migration plan. While existing functionality for service templates will continue to be available in OcNOS 6.6.1, the related commands and configuration procedures have been removed from the documentation to encourage migration to sub-interface-based service binding. This change aligns with our intent of providing a more manageable and operationally usable sub-interface-based configuration model to our OcNOS customer base.



Note: Although service-template-based functionality continues to be available in 6.6.1, customers are advised to begin transitioning existing deployments to sub-interface-based configurations.

Guidance

When upgrading from an earlier release (prior to 6.6.1), review all configurations that use service-template under `mpls-l2-circuit` or `mpls-vpls`, and reconfigure them using sub-interfaces with appropriate VLAN encapsulation.

Key Capabilities

- Sub-interfaces (e.g., `xe21.1`, `xe21.100`) are created to handle encapsulation (e.g., untagged, outer-vlan).

- The `access-if-vpws` or `access-if-vpls` context is used to bind services directly to the sub-interface.
- This model improves operational clarity and scalability.

Example

VPWS Sub-interface-based Configuration (Recommended):

```
interface xel
  no switchport
  no mpls-l2-circuit <vpws name>
  exit
commit

interface xel.1 switchport
  encapsulation untagged
  access-if-vpws
  mpls-l2-circuit <vpws name> primary
  exit
commit
```

VPWS Service Template-based Configuration:

```
service-template <service template name>
match untagged
!
mpls l2-circuit <vpws name> <virtual circuit ID number> <far end neighbor>
!
interface xel
  switchport
  mpls-l2-circuit <vpws name> service-template <service template name> primary
!
```

VPLS Sub-interface-based Configuration (Recommended):

```
interface xel
  no switchport
  no mpls-vpls <vpls instance name> service-template <service template name>
  exit
Commit

interface xel.1 switchport
  encapsulation untagged
  access-if-vpls
  mpls-vpls <vpls instance name>
  exit
commit
```

VPLS Service Template-based Configuration:

```
service-template <service template name>
!
mpls vpls <vpls instance name> <vpls id>
  signaling ldp
  vpls-type <encapsulation type>
```

```
vpls-peer <far end peer address>
exit-signaling
exit-vpls
!
interface xe1
switchport
mpls-vpls <vpls instance name> service-template <service template name>
exit-if-vpls
!
```



Note: It is highly recommended to migrate your Service Template configurations to sub-interfaces prior to the upgrade. Following the upgrade, the Service Template CLI commands will be hidden and no longer accessible for configuration. Refer to the table below for the configuration changes before and after upgrade.

| | Before Upgrade from 6.6.1 | After Upgrade to 7.0 |
|--------------------------------|--|--|
| Service Template Configuration | <pre>! service-template st1 ! service-template st2 ! interface xe8 mpls-l2-circuit t1 service-template st1 primary ! interface xe5 vc-mode revertive service-template st3 mpls-vpls vpls1 service-template st2 !</pre> | <pre>! interface xe8 ! interface xe5 !</pre> |