Key Features

February 2025

IP Infusion Inc.

3979 Freedom Circle

Suite 900

Santa Clara, California 95054

http://www.ipinfusion.com/

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

support@ipinfusion.com

Contents

Preface

This guide describes how to configure OcNOS.

Audience

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

Conventions

Table P-1 shows the conventions used in this guide.

Table 1: Conventions

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

Related Documentation

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

Feature Availability

The features described in this document that are available depend upon the OcNOS SKU that you purchased. See the *Feature Matrix* for a description of the OcNOS SKUs.

Migration Guide

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

Support

For support-related questions, contact support@ipinfusion.com.

Comments

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

Improved Routing

This section describes the new features for improved network routing introduced in the Release 6.6.0.

CHAPTER 1 BGP Automated Dynamic Route Policy Update

Overview

The Border Gateway Protocol (BGP) is a peer routing protocol, and route maps within BGP play an essential role in filtering and modifying routing information to ensure that only the most optimal routes are advertised. BGP route maps use a combination of prefix lists and sequence numbers, which provide precise control over which IP addresses are allowed or denied. These route maps can be applied in both inbound (Adj-RIB-In) and outbound (Adj-RIB-Out) directions for routing updates to BGP peer neighbors.

In OcNOS 6.5.x or earlier releases, route maps are not automatically applied to a BGP peer neighbor upon updates. Manually executing commands, such as *clear ip bgp A.B.C.D* with `soft in/out`, is required for each address family to apply route map changes. This manual intervention is necessary to reflect the changes in the BGP Routing Information Base (RIB).

To address this limitation, a new CLI bgp auto-policy-soft-reset enable, has been introduced. This functionality automatically applies changes to the BGP RIB when route maps or associated lists (such as prefix-list, distribute-list, as-list, access-list, community-list, extended community list, redistribution settings, and more) are applied or updated to a BGP peer neighbor. This dynamic route policy update eliminates the need for the manual execution of the `clear ip bgp <> soft in/out` command for various address families. This greatly simplifies network operations.

The dynamic seamless BGP route policy update introduces the following two new CLIs:

- bgp auto-policy-soft-reset enable
- debug bgp soft-clear

Note: The bgp auto-policy-soft-reset enable functionality is applicable only when apply or remove or update the following:

- route map
- prefix-list
- distribute-list
- as-list
- access-list
- community-list

Note: This functionality is not supported:

- when the peer does not support route refresh messages.
- when the peer is configured with `bgp soft-reconfig-backup.`
- when the router config disables auto soft reset and the administrator prefers manual resets instead.

Benefits

This feature benefits network administrators by reducing the manual effort involved in BGP route filtering and advertisement updates, improving operational efficiency, and ensuring that routing policies are applied immediately without requiring a BGP session reset in real-time propagation of routing information.

Prerequisites

Make sure the provider edge routers and BGP route reflectors have the following active configurations.

- Stable BGP session should exist between Provider edge nodes.
- Support for IPv4, IPv6, VPNv4, BGP Labeled-Unicast, VPNv6 address families on all nodes.
- Established BGP neighbor peers or peer groups are up.
- Ensure baseline BGP configurations are proper.
- Check advertised and received routes for accuracy.

Configuring BGP auto-policy-soft-reset

This procedure outlines the steps to configure the bgp auto-policy-soft-reset enable functionality, which ensures routing policy changes are applied without disrupting active BGP sessions.

Topology

The configuration uses the below topology, which involves two Provider Edge routers PE1, PE2 and an iBGP router act as a Route Reflector (RR) between PE1 and PE2.

Topology of BGP Auto Policy Soft Reset

Note: Before configuration meet all Prerequisites

Procedure

This procedure uses the following interfaces and IP addresses:

- PE1 - `xe3` interface IP address 10.11.1.1/29, loopback IP - 1.1.1.1/32
- PE2 - `ge14` IP address 10.12.3.2/29, loopback IP - 3.3.3.3/32
- RR - `ce46/1` IP address 10.11.1.2/29, loopback IP - 11.11.11.11/32

1. Verify the base BGP configuration on PE1, PE2, and RR before enabling the `bgp soft-clear` CLI. Refer to the sample out shown in PE1 - Before enabling the bgp auto-policy-soft-reset CLI, RR - Before enabling the bgp auto-policy-soft-reset CLI, and PE2- Before enabling the bgp auto-policy-soft-reset CLI.

2. Verify the interface, label-switching, IGP, LDP, existing route maps filter criteria such as prefix-list on PE1 and PE2. Here, is the sample configuration.
   PE1

```
debug bgp soft-clear
.
.
ip vrf management
!
ip vrf vrf221
```

```
 rd 221:1
 route-target both 221:1
!
!
router ldp
 session-protection
 targeted-peer ipv4 3.3.3.3
  exit-targeted-peer-mode
 transport-address ipv4 1.1.1.1
!

interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
!
interface xe3
 description connected-to-RR
 load-interval 30
 ip address 10.11.1.1/29
 ipv6 address 3001::2/64
 mtu 9216
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
router ospf 100
 ospf router-id 1.1.1.1
 network 1.1.1.1/32 area 0.0.0.0
 network 10.2.1.0/29 area 0.0.0.0
 network 10.11.1.0/29 area 0.0.0.0
!
router ipv6 ospf
 router-id 1.1.1.1
!
```

PE2

```
 debug bgp soft-clear
ip vrf vrf221
 rd 221:1
 route-target both 221:1
!
router ldp
 auto-targeted-session
 session-protection
 transport-address ipv4 3.3.3.3
!
interface lo
 ip address 127.0.0.1/8
 ip address 3.3.3.3/32 secondary
```

```
 ipv6 address ::1/128
!
interface ge14
 description connected-to-RR
 load-interval 30
 ip address 10.12.3.2/29
 ipv6 address 6001::3/64
 mtu 9216
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
router ospf 100
 network 3.3.3.3/32 area 0.0.0.0
 network 10.12.3.0/29 area 0.0.0.0
 network 10.12.4.0/29 area 0.0.0.0
!
router ipv6 ospf
 router-id 3.3.3.3
!
```

RR

```
debug bgp soft-clear
!
ip vrf management
!
router ldp
!
interface ce46/1
 description connected-to-PE1
 load-interval 30
 ip address 10.11.1.2/29
 ipv6 address 3001::3/64
 mtu 9216
 label-switching
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
interface lo
 ip address 127.0.0.1/8
 ip address 11.11.11.11/32 secondary
 ipv6 address ::1/128
!
interface xe3
 description connected-to-PE2
 speed 1g
 load-interval 30
 ip address 10.12.3.3/29
 ipv6 address 6001::2/64
 mtu 9216
 label-switching
```

```
 ipv6 router ospf area 0.0.0.0 instance-id 0
 enable-ldp ipv4
!
 exit
!
router ospf 100
 ospf router-id 11.11.11.11
 network 10.11.1.0/29 area 0.0.0.0
 network 10.11.12.0/29 area 0.0.0.0
 network 10.12.3.0/29 area 0.0.0.0
 network 11.11.11.11/32 area 0.0.0.0
!
router ipv6 ospf
 router-id 11.11.11.11
!
```

3. Enable `bgp auto-policy-soft-reset` on PE1, PE2, and RR. Here, is the sample configuration.

PE1

```
!
router bgp 100
 bgp router-id 1.1.1.1
 bgp auto-policy-soft-reset enable
 !
```

PE2

```
!
router bgp 100
 bgp router-id 3.3.3.3
 bgp auto-policy-soft-reset enable
 !
```

RR

```
!
router bgp 100
 bgp router-id 11.11.11.11
 bgp auto-policy-soft-reset enable
 no bgp inbound-route-filter
 neighbor peer_group1 peer-group
 neighbor peer_group1 remote-as 100
 neighbor peer_group1 update-source lo
 neighbor peer_group1 advertisement-interval 0
 neighbor 1.1.1.1 peer-group peer_group1
 neighbor 3.3.3.3 peer-group peer_group1
 !
```

4. Add a route map policy on PE1. For example, configure route map to allow or deny additional prefixes and change the attributes values.

```
!
ip prefix-list prefix1
 seq 5 permit 31.0.1.0/24
!
```

```
 route-map rm_pe1 permit 1
 match ip address prefix-list prefix1
 set metric 200
 set weight 200
 set local-preference 200
 set community 100:101
!
```

5. Add a route map policy on PE2. For example, configure route map to allow or deny additional prefixes and change the attributes values.

```
!
ip community-list standard Community1 permit 100:101
!
 route-map rm_pe2 permit 1
 match community Community1
 set community 100:102
!
```

6. Attach the updated route map filter to BGP peers on PE1 and PE2. For example, on PE1, attach the route map to filter specific routes on outgoing routes using "`neighbor 11.11.11.11 route-map rm-pe1 out`" command towards PE2 under address family vpnv4 unicast. On, PE2, attach it to filter specific routes on incoming routes using "`neighbor 11.11.11.11 route-map rm-pe2 in`" command under address family vpnv4 unicast.

PE1

```
!
router bgp 100
 bgp router-id 1.1.1.1
!
address-family vpnv4 unicast
 neighbor 11.11.11.11 activate
 neighbor 11.11.11.11 route-map rm_pe1 out
 exit-address-family
|
```

PE2

```
!
router bgp 100
bgp router-id 3.3.3.3
!
address-family vpnv4 unicast
 neighbor 11.11.11.11 activate
 neighbor 11.11.11.11 route-map rm_pe2 in
 exit-address-family
!
```

7. Verify the advertised and received routes to ensure the policy changes have taken effect on PE1, PE2 and RR. Refer to the sample out shown in PE2 - After enabling the bgp auto-policy-soft-reset.

Validation

PE1 - Before enabling the bgp auto-policy-soft-reset CLI

The following show output displays the active BGP configuration PE1

```
PE1#show running-config bgp
!
router bgp 100
 bgp router-id 1.1.1.1
 neighbor 11.11.11.11 remote-as 100
 neighbor 11.11.11.11 update-source lo
 neighbor 11.11.11.11 advertisement-interval 0
 !
 address-family ipv4 unicast
 network 1.1.1.1/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family ipv6 unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family ipv4 vrf vrf221
 redistribute connected
 redistribute ospf
 neighbor 101.1.1.10 remote-as 200
 neighbor 101.1.1.10 activate
 exit-address-family
 !
 address-family ipv6 vrf vrf221
 redistribute connected
 exit-address-family
 !
 exit
!
PE1#
PE1#show ip bgp labeled-unicast summary
BGP router identifier 1.1.1.1, local AS number 100
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
11.11.11.11       4   100         49         36       1      0      0  00:11:55             0
```

```
Total number of neighbors 1

Total number of Established sessions 1
PE1#
PE1#show ip bgp vpnv4 vrf vrf221
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

     Network          Next Hop            Metric    LocPrf     Weight Path
Route Distinguisher: 221:1 (Default for VRF vrf221)
*> l 31.0.0.0/24      101.1.1.10            0         100        0   200 i
*> l 31.0.1.0/24      101.1.1.10            0         100        0   200 i
*> l 31.0.2.0/24      101.1.1.10            0         100        0   200 i
*> l 33.0.0.0/24      101.1.1.10            0         100        0   200 i
*> l 33.0.1.0/24      101.1.1.10            0         100        0   200 i
*> l 33.0.2.0/24      101.1.1.10            0         100        0   200 i
*>il 33.0.3.0/24      3.3.3.3               0         100        0   200 i
*>il 33.0.4.0/24      3.3.3.3               0         100        0   200 i
*> l 101.1.1.0/24     0.0.0.0               0         100        32768  ?
*>il 151.1.1.0/24     3.3.3.3               0         100        0    ?
 Announced routes count = 7
 Accepted routes count = 3
PE1#
PE1#show ip bgp vpnv4 vrf vrf221 31.0.1.0
Route Distinguisher: 221:1 (Default for VRF vrf221) Routing Entry for prefix: 31.0.1.0/
24
  Advertised to non peer-group peers:
  11.11.11.11 11.11.11.11
  AS path:200
  Path Selection reason: Nothing left to compare
   Nexthop:101.1.1.10 (IGP metric 0) from 101.1.1.10 (Remote Id:192.0.4.53) Peer
nexthop: 101.1.1.10
      Origin IGP, metric 0, localpref 100, Out-label 0, In-label 24384, refcnt: 1
      valid, external, best, source-safi: 1
      Extended Community: RT:221:1
      rx path_id: -1      tx path_id: -1
      Add-Path Announcement: Not advertised to any peer
      Last update: Fri Dec  6 18:11:29 2024, 00:12:28 ago
```

RR - Before enabling the bgp auto-policy-soft-reset CLI

The following show output displays the active BGP configuration on RR b

```
RR#show running-config bgp
!
router bgp 100
 bgp router-id 11.11.11.11
 bgp auto-policy-soft-reset enable
 no bgp inbound-route-filter
 neighbor peer_group1 peer-group
 neighbor peer_group1 remote-as 100
```

```
 neighbor peer_group1 update-source lo
 neighbor peer_group1 advertisement-interval 0
 neighbor 1.1.1.1 peer-group peer_group1
 neighbor 3.3.3.3 peer-group peer_group1
 !
 address-family ipv4 unicast
 network 11.11.11.11/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 neighbor peer_group1 activate
 neighbor peer_group1 route-reflector-client
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor peer_group1 activate
 neighbor peer_group1 route-reflector-client
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor peer_group1 activate
 neighbor peer_group1 route-reflector-client
 exit-address-family
 !
 address-family ipv6 unicast
 neighbor peer_group1 activate
 neighbor peer_group1 route-reflector-client
 exit-address-family
 !
 exit
!
RR#
RR#show ip bgp labeled-unicast summary
BGP router identifier 11.11.11.11, local AS number 100
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
1.1.1.1           4   100        124        142       1      0      0  00:09:33            0
3.3.3.3           4   100        140        131       1      0      0  00:01:14            0

Total number of neighbors 2

Total number of Established sessions 2
RR#
RR#show ip bgp vpnv4 all summary
BGP router identifier 11.11.11.11, local AS number 100
BGP table version is 14
2 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        124        142      14      0      0  00:09:47             7
3.3.3.3           4   100        140        131      14      0      0  00:01:28             3

Total number of neighbors 2

Total number of Established sessions 2
RR#
```

PE2- Before enabling the bgp auto-policy-soft-reset CLI

The following show output displays the active BGP configuration on PE2.

```
PE2#
PE2#show running bgp
!
router bgp 100
 bgp router-id 3.3.3.3
 bgp auto-policy-soft-reset enable
 neighbor 11.11.11.11 remote-as 100
 neighbor 11.11.11.11 update-source lo
 neighbor 11.11.11.11 advertisement-interval 0
 !
 address-family ipv4 unicast
 network 3.3.3.3/32
 exit-address-family
 !
 address-family ipv4 labeled-unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 11.11.11.11 activate
 neighbor 11.11.11.11 route-map rm_pe2 in
 exit-address-family
 !
 address-family vpnv6 unicast
 neighbor 11.11.11.11 activate
 exit-address-family
 !
 address-family ipv6 unicast
 exit-address-family
 !
 address-family ipv4 vrf vrf221
 redistribute connected
 redistribute ospf
 neighbor 151.1.1.10 remote-as 200
 neighbor 151.1.1.10 activate
 exit-address-family
 !
 address-family ipv6 vrf vrf221
 redistribute connected
```

```
 exit-address-family
 !
 exit
!
PE2#
PE2#show ip bgp labeled-unicast summary
BGP router identifier 3.3.3.3, local AS number 100
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
11.11.11.11       4   100         45         55       1      0      0  00:05:17            0

Total number of neighbors 1

Total number of Established sessions 1

PE2#show ip bgp vpnv4 vrf vrf221
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

     Network          Next Hop            Metric    LocPrf     Weight Path
Route Distinguisher: 221:1 (Default for VRF vrf221)
*>il 31.0.0.0/24      1.1.1.1               0         100        0   200 i
*>il 31.0.1.0/24      1.1.1.1               0         100        0   200 i
*>il 31.0.2.0/24      1.1.1.1               0         100        0   200 i
*>il 33.0.0.0/24      1.1.1.1               0         100        0   200 i
*>il 33.0.1.0/24      1.1.1.1               0         100        0   200 i
*>il 33.0.2.0/24      1.1.1.1               0         100        0   200 i
*> l 33.0.3.0/24      151.1.1.10            0         100        0   200 i
*> l 33.0.4.0/24      151.1.1.10            0         100        0   200 i
*>il 101.1.1.0/24     1.1.1.1               0         100        0    ?
*> l 151.1.1.0/24     0.0.0.0               0         100        32768  ?
 Announced routes count = 3
 Accepted routes count = 7
PE2#
PE2#show ip bgp vpnv4 vrf vrf221 31.0.1.0
Route Distinguisher: 221:1 (Default for VRF vrf221) Routing Entry for prefix: 31.0.1.0/
24
  Not advertised to any peer
  AS path:200
  Path Selection reason: Nothing left to compare
  Nexthop:1.1.1.1 (IGP metric 3) from 11.11.11.11 (Originator Id:1.1.1.1) (Remote Id:11.11.11.11) Peer nexthop: 11.11.11.11
     Origin IGP, metric 0, localpref 100, Out-label 24384      valid, internal, best, source-safi: 128
       Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External, imported)
       Extended Community: RT:221:1
       Originator: 1.1.1.1, Cluster list: 11.11.11.11
       rx path_id: -1      tx path_id: -1
       Add-Path Announcement: Not advertised to any peer
```

```
      Last update: Fri Dec  6 18:19:48 2024, 00:05:25 ago
```

PE1 - After enabling the bgp auto-policy-soft-reset

Following show output displays the modified route policy.

```
PE1#show running-config route-map
!
route-map rm_pe1 permit 1
 match ip address prefix-list prefix1
 set metric 200
 set weight 200
 set local-preference 200
 set community 100:101
!
PE1#
```

PE2 - After enabling the bgp auto-policy-soft-reset

The following show output displays the active BGP configuration on PE2.

```
PE2#show ip bgp vpnv4 vrf vrf221 31.0.1.0
Route Distinguisher: 221:1 (Default for VRF vrf221) Routing Entry for prefix: 31.0.1.0/
24
  Not advertised to any peer
  AS path:200
  Path Selection reason: Nothing left to compare
  Nexthop:1.1.1.1 (IGP metric 3) from 11.11.11.11 (Originator Id:1.1.1.1) (Remote Id:11.11.11.11) Peer nexthop: 11.11.11.11
     Origin IGP, metric 200, localpref 200, Out-label 24384      valid, internal, best, source-safi: 128
     Duplicated: (source VRF-ID: 0, source VRF: DEFAULT, VRF-External, imported)
     Community: 100:102
       Extended Community: RT:221:1
       Originator: 1.1.1.1, Cluster list: 11.11.11.11
       rx path_id: -1      tx path_id: -1
       Add-Path Announcement: Not advertised to any peer
       Last update: Mon Dec  2 23:18:45 2024, 00:17:31 ago

PE2#
```

Implementation Examples

Following are some of the use cases where manual clearance `clear ip bgp A.B.C.D soft out` of route updates are necessary before applying them to a neighbor peer or peer group.

Scenario 1:

When configuring `bgp additional-paths` with options `send|receive|send-receive|select (all | best<2-3>))` or when enabling `neighbor advertise additional-paths`, route change occurs that needs to be updated in the neighbor peer or peer group through BGP RIB-out.

Scenario 2:

Configuring `neighbor unsuppress-map` leads to route changes that needs to be announced in BGP RIB-out.

Scenario 3:

Enabling `neighbor allowas-in` causes change in routes that needs to accepted via BGP RIB-in.

Scenario 4:

When multiple BGP address families facilitate route management between two customer edge devices (CE1 and CE2) via Route Reflector, enabling bgp auto-policy-soft-reset enable dynamically updates newly announced route map prefixes.

CLI Commands

Following are the new CLIs introduced to clear the route-map updates automatically for each BGP address-family.

bgp auto-policy-soft-reset enable

Use this command ONLY to enable BGP soft clear for address families IPv4, IPv6, VPNv4, VPNv6, labeled-unicast regarding route map, table-map, prefix-list, distribute-list, as-list, access-list or community-list updates to apply the modified routes and attributes automatically on the neighbor peer or peer group.

Use the `no` form of this CLI to disable the soft reset.

Note: EVPN address family is not supported.

Command Syntax

```
bgp auto-policy-soft-reset enable
no bgp auto-policy-soft-reset enable
```

Parameters

None

Default

Disabled

Command Mode

BGP Router Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Entering the CLI enables the soft clear the route map.

```
OcNOS(config)#router bgp
OcNOS(config-router)#bgp auto-policy-soft-reset enable
OcNOS(config-router)#address-family vpnv4 unicast
OcNOS(config-router-af)#neighbor 50.50.50.50 prefix-list prefix-inside in
OcNOS(config-router-af)#commit
OcNOS(config-router-af)#exit
OcNOS(config-router)#commit
```

debug bgp soft-clear

Use this command to log the BGP auto soft clear information with timestamp for debug purpose.

Command Syntax

```
debug bgp soft-clear
```

Parameters

None

Default

Disabled

Command Mode

Configure Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Entering the CLI enables the BGP soft clear log file.

```
OcNOS(config)#debug bgp soft-clear
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
AS	Autonomous System
BGP	Border Gateway Protocol router
PE	Provider Edge router
RIB	Route Information Base
RR	Route-Reflector router
Route-map	Maps used for filtering RIB in/out on various use cases
Prefix list	List created on router with Network prefixes *ge*, *le* options too to indicate network address length
Access List	List created on a router with IP addresses mentioned (IPv4, IPv6) without any network address length flexibility like prefix lists.
AS List	Autonomous System path list where BGP can filter routes based on AS path it checks in the control packet
Community List	Community is a mini group inside AS, and BGP can filter these communities based on AS:NN where the first 16 bits are AS, next 16 bits a unique number with AS.

CHAPTER 2 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 failsafe 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:

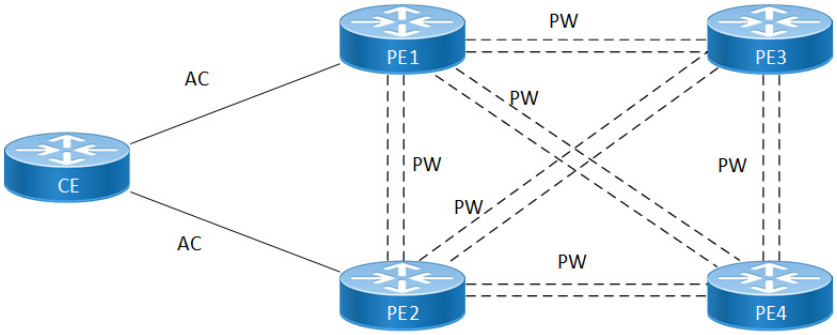

Feature Characteristics

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.

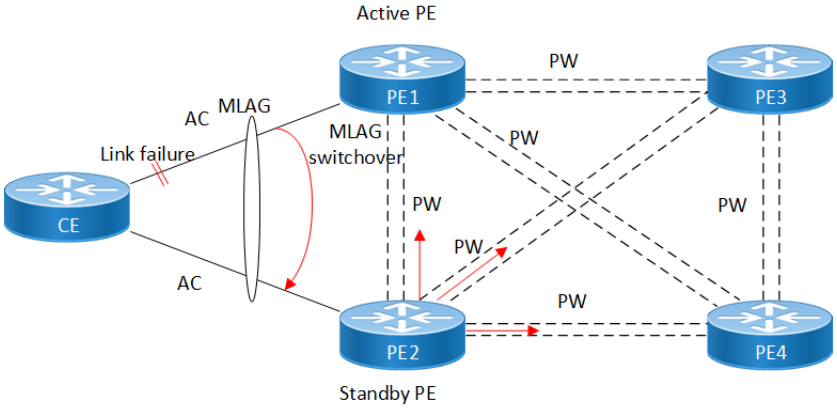

The above figure shows a VPLS mesh linked to a dual-homed CE connected to PE1 and PE2. The traffic flows from CE to PE3 through the active MLAG PE1, while PE2 remains a standby MLAG. When PE1 experiences a link failure, the standby MLAG becomes active, and traffic flows through the PE2 device to reach 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.

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 2-1: MLAG VPLS configuration topology

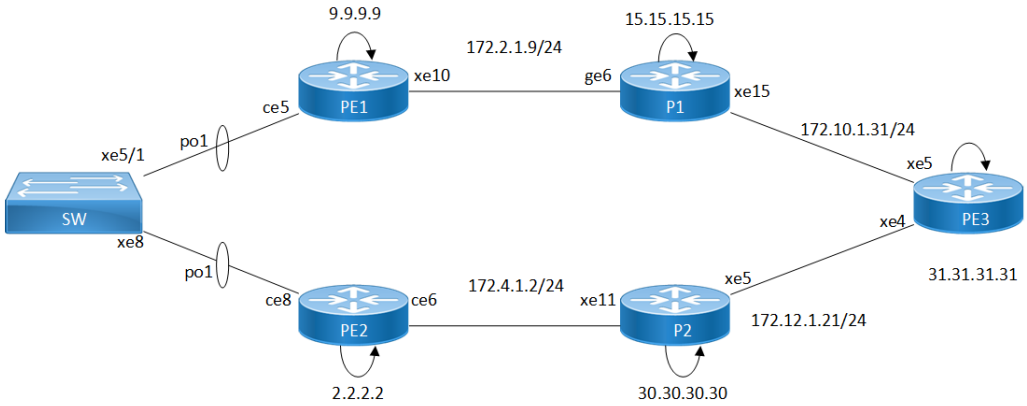

MLAG VPLS Configuration

The following steps show the configuration of the MLAG VPLS with automatic MAC flush capability during the MLAG switchover.

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 po1.200 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 po1
 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 po1.200 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 po1
 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 po1.200 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
2/Up                    2         LDP             Active     00:20:29      -

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
0/Dn                    4         LDP              Standby     -   -

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    po1.200

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 Acive MLAG peer interface on PE1 or on switch)

```
PE1(config)#interface po1
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         LDP           Active     00:21:54      -
200        31.31.31.31       24325         25602      ce6            26890      2/Up
4         LDP           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         LDP           Active     17:45:58   -
200        9.9.9.9           N/A           26885      N/A            25604      0/Dn
2         LDP           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.

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
```

Revised CLI Commands

The following command is revised:

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*.

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.
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).

CHAPTER 3 Flex Algorithm for ISIS

Overview

Flexible Algorithms (Flex-Algo) is a cutting-edge enhancement within Segment Routing (SR) technology, enabling customize path computation within IGP protocols like IS-IS. This feature supports diverse traffic engineering needs, such as low latency or bandwidth optimization, without external controllers. Flex-Algo creates logical routing planes within an IGP, enabling constrained traffic paths tailored to application-specific requirements. This makes it a key component in networks supporting modern services, including 5G.

Note: Currently, OcNOS supports this feature for IS-IS as the IGP.

The Need for ISIS Flexible Algorithms

Traditional IGPs, using Shortest Path First (SPF) computations, limit flexibility by routing traffic solely based on link costs. This results in static traffic patterns, leading to underutilization of network resources and an inability to meet specific requirements, such as latency minimization or link exclusion.

While solutions like MPLS-TE addressed these challenges, they introduced complexity, scalability issues, and increased operational overhead. Flex-Algo simplifies traffic engineering by allowing operators to define routing rules directly within the IGP. This enables efficient handling of traffic for diverse applications, particularly in the 5G era, where flexibility is crucial.

Feature Characteristics

Flex-Algo Definition

Flex-Algo allows the creation of up to 128 unique algorithms, each operating independently. Key parameters defining a Flex-Algo include:

Path Computation Method:

- **Standard SPF:** Uses the Dijkstra algorithm for shortest path computation and allows policy-driven modifications.
- **Strict SPF:** Similar to Standard SPF but restricts policy-based adjustments.

Routing Metrics:

- IGP Metric: Based on traditional link costs as defined in the IGP.
- Traffic Engineering (TE) Metric: Optimizes paths using TE attributes like bandwidth and utilization.
- Link Delay: Routes traffic based on the lowest unidirectional delay.

Priority Levels:

- Algorithms are processed based on their priority value.
- In cases of identical priority, the System-ID acts as a tiebreaker, with the advertisement having the highest System-ID selected.

Link Affinity Constraints:

Defines constraints for path computation, using affinity attributes configured as Administrative Groups (AG) or Extended Administrative Groups (EAG):

- **Forward EAG Exclude Any:** Excludes links with any matching affinity bits.

- **Forward EAG Include Any:** Includes links with at least one matching affinity bit.
- **Forward EAG Include All:** Includes links only if all specified affinity bits match.
- **Reverse EAG Exclude Any:** Excludes links from reverse path computation if any matching affinity bits are found.
- **Reverse EAG Include Any:** Includes links if at least one affinity bit matches.
- **Reverse EAG Include All:** Includes links only if all affinity bits match.

Delay Constraints:

Routes traffic over low-delay links using metrics collected via protocols like TWAMP, optimizing latency-sensitive applications.

Participation in a Flexible Algorithm

Flex-Algo Support Advertisement

Routers advertise support for specific Flex-Algos using IGP. Algorithm values (128–255) are tightly coupled with Prefix-SIDs, enabling algorithm-specific forwarding. Operators configure routers to participate in algorithms based on network requirements.

Flex-Algo Definition Advertisement

To ensure loop-free forwarding, routers in the network must share a consistent algorithm definition. Routers advertise these definitions, including metrics and affinity constraints, using a priority mechanism. Recommendations include:

At least one router per area must advertise the algorithm definition. Configuring two routers ensures redundancy.

Without a valid algorithm definition advertisement, Flex-Algo cannot function.

Flex-Algo Prefix-SID Advertisement

Routers participating in a Flex-Algo advertise MPLS-labeled paths associated with the algorithm's Prefix-SID. Prefix-SIDs are specific to the algorithm and enable algorithm-driven forwarding. Rules include:

- Only prefixes associated with an algorithm-specific Prefix-SID are included in the forwarding table.
- Prefix-SIDs and prefixes can be leaked between areas but are limited to reachable Layer 1 or Layer 2 routers.

Path Calculation and Forwarding Entries

Path Calculation:

- Nodes not supporting the algorithm are excluded.
- Links with excluded affinities are pruned from the topology.
- Only links advertising the algorithm’s required metric are considered.

Forwarding Entries:

- Paths must be installed using the Prefix-SID advertised for the algorithm.
- If no Prefix-SID is available, the path is not installed in the forwarding table.
- Routes for each algorithm are installed in algorithm-specific RIB groups. By default, IS-IS Flex-Algo routes are added to MPLS RIBs.

Example of Flex-Algo Path Calculation

Nodes and Links:

- The topology consists of 6 nodes (1, 2, 3, 4, 5, and 6).
- Each link is assigned:
  - Cost (IGP Metric): Default is 10.
  - Delay (TE Metric): Some links (for example: Node 1 to Node 5 and Node 5 to Node 6) have a delay of 6, while others have a delay of 10.

Flex-Algo 0 (Default Algorithm)

- Uses cost (IGP metric) as the path computation metric.
- Considers all links in the network (no exclusions based on affinities).
- Equal-cost paths are possible due to the uniform cost across links.

Flex-Algo 128 (Custom Algorithm)

- Uses delay (TE metric) for path computation.
- Only includes links marked with the green affinity attribute.
- Excludes all other links from the topology unless they meet the "green" affinity condition.

Overview of Both Algorithms

- Shows how both Flex-Algo 0 (blue) and Flex-Algo 128 (green) are configured within the same topology.
- Links supporting Flex-Algo 128 are marked in green and constrained by the green affinity attribute.

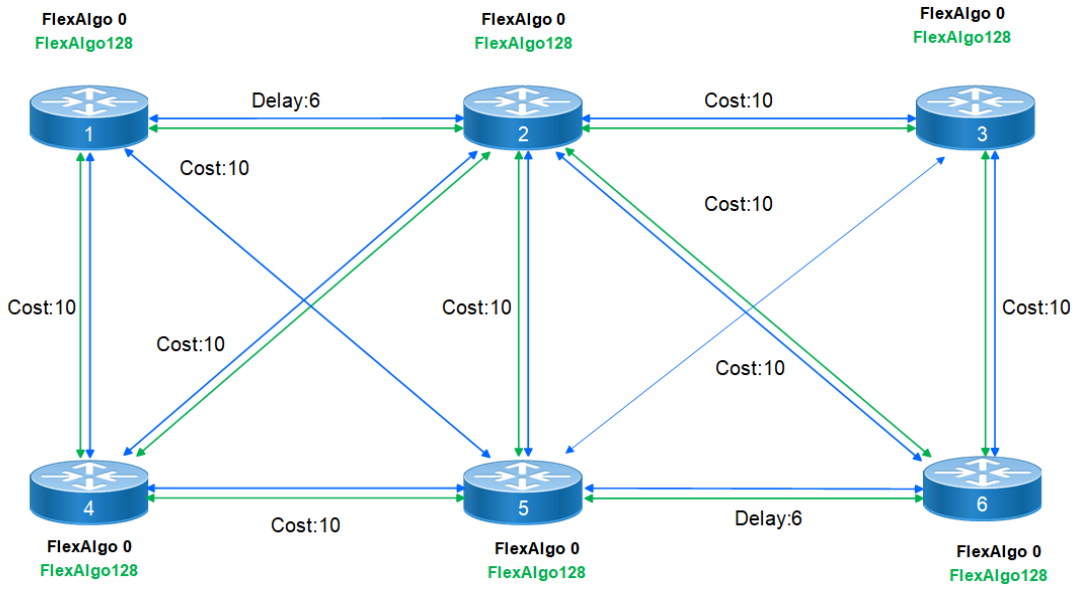

Flex-Algo 0 Topology

- All links are included in this topology, as Flex-Algo 0 uses IGP metrics and has no constraints.
- Path Example: For traffic from Node 1 to Node 3, the paths are:
  - Node 1 → Node 2 → Node 3.
  - Node 1 → Node 5 → Node 3.
- Both paths are valid since the total cost for each path is equal (20).

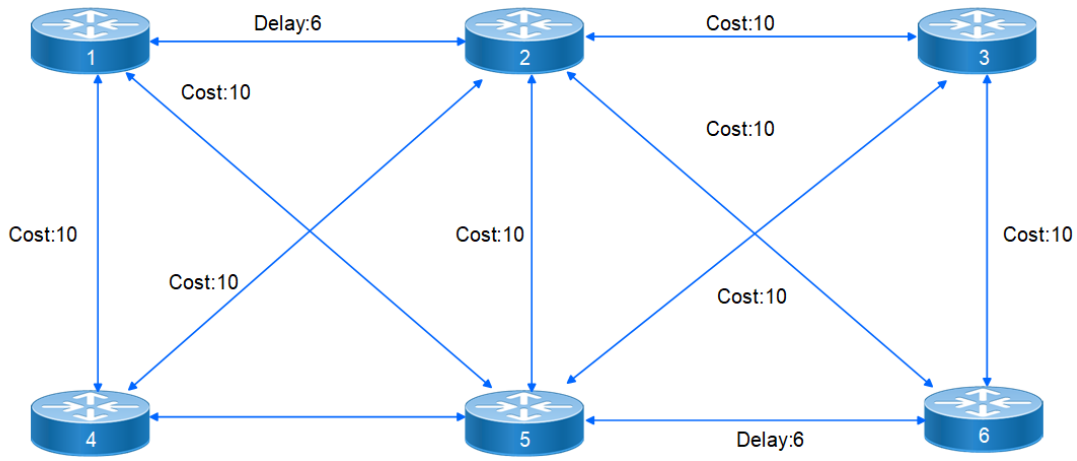

Flex-Algo 128 Topology

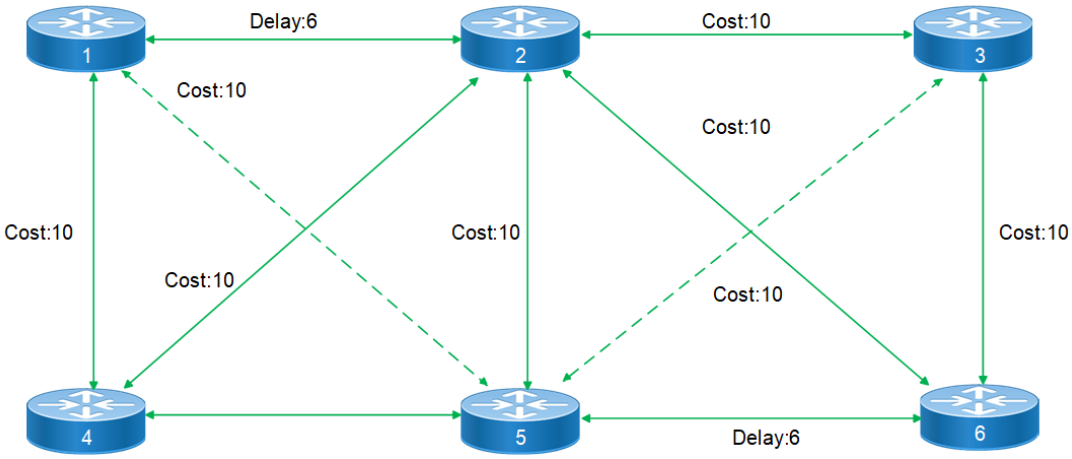

- Only includes links with the green affinity attribute.
- Some links (for example: Node 1 → Node 4, Node 4 → Node 5) are excluded because they do not have the required affinity.
- Path Example: For traffic from Node 1 to Node 3, the only valid path is:
  - Node 1 → Node 2 → Node 3.
  - This path is selected because it minimizes delay (TE metric) while satisfying the "green" affinity constraint.

Key Observations

Flex-Algo 0:

- Suitable for general-purpose routing without any specific constraints.
- Uses cost as the metric, enabling equal-cost path computation.

Flex-Algo 128:

- Optimized for delay-sensitive traffic, focusing on paths with minimal delay.
- Excludes links without the green affinity, reducing the available topology and forcing path selection based on both constraints and delay.

Benefits

Flex-Algo introduces significant advantages for network operations:

- **Tailored Path Selection:** Enables customization of routing paths to meet specific traffic and application requirements, such as low latency or high throughput.
- **Integrated Traffic Engineering:** Embeds traffic engineering capabilities within IGPs, reducing reliance on external controllers.
- **Optimized Network Resources:** Improves utilization by avoiding congestion and ensuring balanced traffic distribution.
- **Simplified Operations:** Reduces the complexity of traditional configurations, such as MPLS-TE, while maintaining scalability for modern networks.

Prerequisites

- Devices must support IS-IS with Segment Routing and Flex-Algo capabilities.
- IGP and SR configurations must be enabled on all participating routers.
- Prefix-SIDs must be assigned for each Flex-Algo.
- Affinity groups (AG or EAG) should be pre-configured to define link constraints.
- Traffic Engineering and SR configurations must be enabled to advertise Sub-TLVs.

Configuration

The following configuration enables Flex Algo with ISIS-SR.

Topology

This topology represents the application of Flex-Algo in a SR network environment, with three distinct algorithms highlighted: Flex Algo 0, 128, and 129. Each algorithm defines a specific logical topology based on unique constraints and use cases.

Topology Visualization:

- Red Lines (Flex Algo 0): Represents the default IGP metric-based paths where all links contribute to general traffic forwarding.
- Green Lines (Flex Algo 128): Highlights paths optimized for bandwidth efficiency using GREEN-affinity links, filtering out non-compliant links.
- Blue Lines (Flex Algo 129): Displays paths optimized for delay-sensitive traffic using BLUE-affinity links, excluding all others.

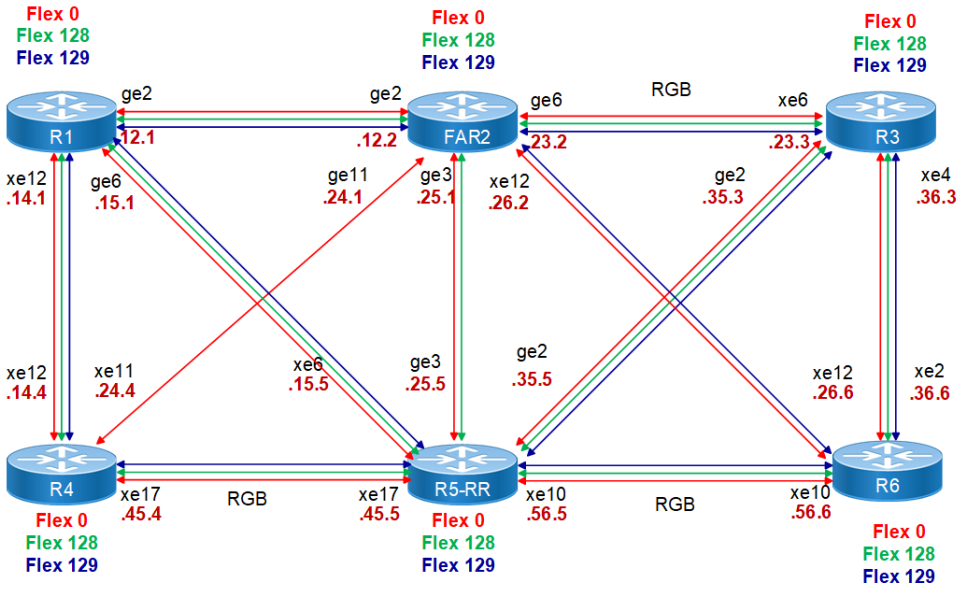

Logical Topologies Based on Flexible Algorithms

Default IGP Metric-Based Topology (FlexAlgo 0)

This topology uses the standard IGP (Interior Gateway Protocol) metrics for shortest-path computation.

Metrics for links default to 10, unless explicitly modified.

Logical Representation: Includes all routers and links in the network. Designed to provide a generic network topology without special constraints or optimizations.

Use Case: Suitable for non-critical, general-purpose traffic. Lacks any special optimization for delay, bandwidth, or other constraints.

Example: Basic network connectivity for general applications like email or web browsing, where performance is not a critical factor.

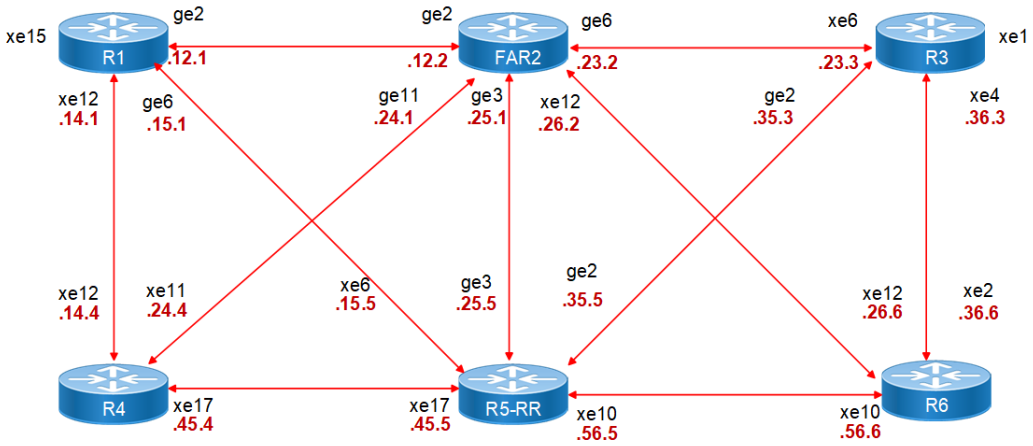

Default IGP Metric-Based Topology (FlexAlgo 0)

TE Metric-Optimized Topology (Flex Algo 128)

Focuses on Traffic Engineering (TE) metrics to optimize network paths. Considers all links with configured TE metrics, such as bandwidth utilization.

Logical Representation: Includes only links with a specific Explicit Affinity Group (EAG) marked as GREEN. Other links (for example: GREY) are excluded, ensuring that the topology adheres to specific affinity constraints.

Use Case: Ideal for scenarios requiring bandwidth efficiency or load balancing. Used for traffic engineering when path optimization is critical.

Example: Managing high-throughput services like data center interconnections or streaming services. Only GREEN-affinity links are used to meet bandwidth requirements, while other links are excluded.

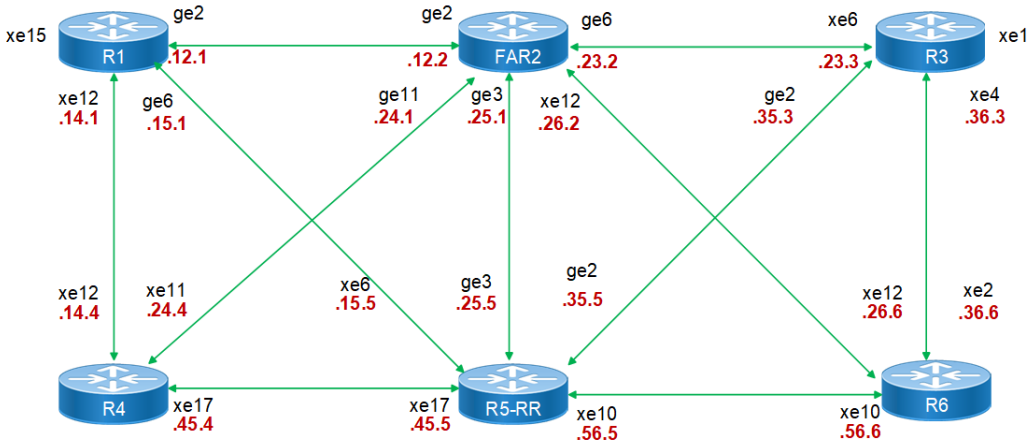

TE Metric-Optimized Topology (FlexAlgo 128)

Delay Metric-Optimized Topology (Flex Algo 129)

A topology designed to minimize delay metrics across the network. Delay is measured either dynamically or configured statically to ensure reliability and predictability.

Logical Representation: Includes links with EAG BLUE to ensure low-delay paths.

Links not marked BLUE are filtered out to exclude high-latency links.

Use Case: Specifically suited for delay-sensitive traffic such as:

- Voice-over-IP (VoIP)
- Live streaming
- 5G applications requiring minimal latency.

Example: Using BLUE-affinity links exclusively for critical traffic, such as real-time communication or interactive applications.

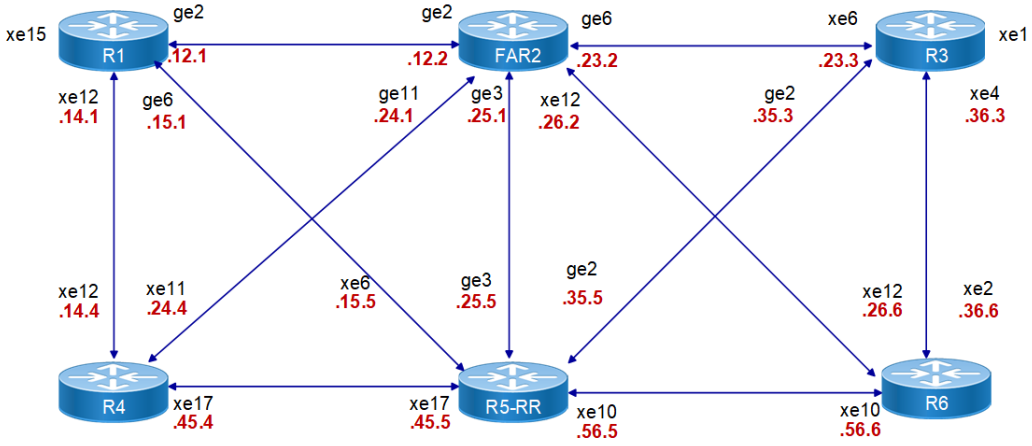

Delay Metric-Optimized Topology (FlexAlgo 129

To configure Flex Algo functionality on nodes with ISIS as IGP, follow the steps mentioned below:

1. Configure loop-back interfaces.
    1. Access interface configuration mode for the loopback interface (`interface loopback1`).
    2. Assign an IPv4 address to the loopback interface using the IPv4 address command followed by the desired IPv4 address and subnet mask (`ip address 25.0.0.1/32`).
    3. Assign appropriate prefix-sid index for the loopback interface (`prefix-sid index 1 no-php`).
    4. Assign appropriate prefix-sid index for flex algorithm for the same loopback interface
    5. Configure IS-IS for IPv4 on the loopback interface using the ip router isis command, specifying the IS-IS process ID (`ip router isis 1`).

```
R1(config)#interface loopback2
R1(config-if)# ip address 25.0.1.1/32 secondary
R1(config-if)# prefix-sid index 1 no-php
R1(config-if)# prefix-sid algorithm-num 128 index 1281
R1(config-if)# prefix-sid algorithm-num 129 index 1291
R1(config-if)# ip router isis 1
R1(config-if)# exit
```

2. Configure network interface.
    1. Access interface configuration mode for the desired network interface (`interface ge2 and ge6 and ge12`).
    2. Assign an IPv4 address to the loopback interface using the ipv4 address command followed by the desired IPv4 address and subnet mask (`ip address 11.0.12.2/24`).
    3. Configure the MTU for the interface (`mtu 9216`).
    4. Configure IS-IS for IPv4 on the interface using the IP router ISIS command, specifying the IS-IS process ID (`ip router isis 1`).

```
R1(config)#interface ge2
R1(config-if)# 11.0.12.2/24
R1(config-if)# mtu 9216
R1(config-if)# label-switching
```

```
R1(config-if)# ip router isis 1
R1(config-if)#exit
```

3. IS-IS & SR configuration:
    1. Set IS-IS metrics (default is 10).
    2. Configure IS-IS for Level 2 with wide metrics.
    3. Enable Segment Routing on all routers:
    4. Configure Node-SID for each router and advertise it:
    5. Configure the New SRGB/SRLB range:

```
R1(config-router)# isis 1
R1(config-router)# node-sid index 1
R1(config)# segment-routing
R1(config-sr)# global block 20000 80000
```

    Here details on IS-IS parameters:
    - Remaining LSP lifetime = 65535 seconds
    - LSP refresh (if configurable) = 65000 seconds
    - LSP ignore error (ISO:2002 RFC 3719)
4. Flex Algo Configuration:
    1. Enable Flex Algo routing on each router:

```
R1(config-router)# capability flex-algo routing
```

    2. Configure Flex Algo 128 and 129 with specific metric types:
        - Configure Flex Algo 128 to minimize TE metrics and include all links.

```
R1(config-router)# flex-algo 128
R1(config-isis-fa)# metric-type te-metric
R1(config-isis-fa)# exit-flex-algo
```

        - Configure Flex Algo 129 to optimize for minimizing delay metrics to ensure predictability and includes all links.

```
R1(config-router)# flex-algo 129
R1(config-isis-fa)# metric-type link-delay
R1(config-isis-fa)# exit-flex-algo
```

5. Configure link attributes for Flex Algo.
    1. Assign admin-group and extended admin-group attributes:
        - **Admin-Group Red**: Assigned to bit position 1.
        - **Extended Admin-Group Green**: Assigned to bit position 32
        - **Extended Admin-Group Blue**: Assigned to bit position 93

```
R1(config)# admin-group red 1
R1(config)# extended-admin-group green 32
R1(config)# extended-admin-group blue 93
```

    2. Apply Link Attributes to Interfaces
    3. Access the desired interface configuration mode
    4. Apply the Admin-Group attribute for Flex Algo:

5. Apply the Extended Admin-Group attributes for Flex Algo:

```
R1(config)# interface ge2
R1(config-if)# isis admin-group flex-algo red
R1(config-if)# isis extended-admin-group flex-algo green
R1(config-if)# isis extended-admin-group flex-algo blue
R1(config-if)# exit
```

6. Define Flex-Algo advertisement.
   - Configure R2 to advertise Flex Algo 128 with a priority of 131:

```
R2(config-router)#flex-algo 128
R2(config-isis-fa)#priority 131
R2(config-isis-fa)#affinity-eag-include-any green
R2(config-isis-fa)#exit-flex-algo
```

   - Configure R5 to serve as the backup node, and advitise the Flex-Algo with a priority of 130.
   - Configure R2 to advertise Flex Algo 129 with a priority of 131:

```
R2(config-router)#flex-algo 129
R2(config-isis-fa)#priority 131
R2(config-isis-fa)#affinity-eag-include-all blue
R2(config-isis-fa)#exit-flex-algo
```

Configuration Sanapshot:

R1

```
qos enable
!
hostname RTR1
!
admin-group red 1
extended-admin-group green 32
extended-admin-group blue 93
!
router-id 25.0.1.1
!
segment-routing
 global block 20000 80000
!
interface ge2
 load-interval 30
 ip address 11.0.12.1/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface ge6
 load-interval 30
 ip address 11.0.15.1/24
 mtu 9216
 label-switching
 isis network point-to-point
```

```
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface loopback1
 ip address 25.0.0.1/32 secondary
 ip router isis 1
!
interface loopback2
 ip address 25.0.1.1/32 secondary
 prefix-sid index 1 no-php
 prefix-sid algorithm-num 128 index 1281
 prefix-sid algorithm-num 129 index 1291
 ip router isis 1
!
interface xe12
 load-interval 30
 ip address 11.0.14.1/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
router isis 1
 is-type level-2-only
 ignore-lsp-errors
 capability flex-algo routing
 flex-algo 129
  metric-type link-delay
 exit-flex-algo
 !
 flex-algo 128
  metric-type te-metric
 exit-flex-algo
 !
 lsp-gen-interval 1
 max-lsp-lifetime 65535
 spf-interval-exp 50 5000
 metric-style wide
 mpls traffic-eng router-id 25.0.1.1
 mpls traffic-eng level-2
 dynamic-hostname
 set-overload-bit on-startup wait-for-bgp
 bfd all-interfaces
 net 49.0000.0100.0000.1011.00
```

```
 isis segment-routing global block 20000 80000
 segment-routing mpls
!
router bgp 65010
 neighbor 25.0.0.5 remote-as 65010
 neighbor 25.0.0.5 update-source loopback1
 neighbor 25.0.1.5 remote-as 65010
 neighbor 25.0.1.5 update-source loopback2
 !
 address-family ipv4 unicast
 neighbor 25.0.0.5 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 25.0.1.5 activate
 exit-address-family
 !
 exit
!
```

R2

```
hardware-profile filter twamp-ipv4 enable
qos enable
!
hostname RTR2
admin-group red 1
extended-admin-group green 32
extended-admin-group blue 93
!
router-id 25.0.1.2
!
segment-routing
 global block 20000 80000
!
interface ge2
 load-interval 30
 ip address 11.0.12.2/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface ge3
 load-interval 30
 ip address 11.0.25.2/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
```

```
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
!
interface ge6
 load-interval 30
 ip address 11.0.23.2/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface ge11
 load-interval 30
 ip address 11.0.24.2/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
!
interface loopback1
 ip address 25.0.0.2/32 secondary
 ip router isis 1
!
interface loopback2
 ip address 25.0.1.2/32 secondary
 prefix-sid index 2 no-php
 prefix-sid algorithm-num 128 index 1282
 prefix-sid algorithm-num 129 index 1292
 ip router isis 1
!
interface xe12
 load-interval 30
 ip address 11.0.26.2/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo blue
!
router isis 1
 is-type level-2-only
```

```
 ignore-lsp-errors
 capability flex-algo routing
 flex-algo 128
  metric-type te-metric
  priority 131
  affinity-eag-include-any green
 exit-flex-algo
 !
 flex-algo 129
  metric-type link-delay
  priority 131
  affinity-eag-include-all blue
 exit-flex-algo
 !
 lsp-gen-interval 1
 max-lsp-lifetime 65535
 spf-interval-exp 50 5000
 metric-style wide
 mpls traffic-eng router-id 25.0.1.2
 mpls traffic-eng level-2
 dynamic-hostname
 set-overload-bit on-startup wait-for-bgp
 bfd all-interfaces
 net 49.0000.0100.0000.1022.00
 redistribute isis level-2 into level-1
 isis segment-routing global block 20000 80000
 segment-routing mpls
```

R3

```
qos enable
!
hostname RTR3
admin-group red 1
extended-admin-group green 32
extended-admin-group blue 93
!
router-id 25.0.1.3
!
segment-routing
 global block 20000 80000
!
interface ge2
 load-interval 30
 ip address 11.0.35.3/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface loopback1
```

```
 ip address 25.0.0.3/32 secondary
 ip router isis 1
!
interface loopback2
 ip address 25.0.1.3/32 secondary
 prefix-sid index 3 no-php
 prefix-sid algorithm-num 128 index 1283
 prefix-sid algorithm-num 129 index 1293
 ip router isis 1
!
interface xe4
 load-interval 30
 ip address 11.0.36.3/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface xe6
 speed 1g
 load-interval 30
 ip address 11.0.23.3/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
router isis 1
 is-type level-2-only
 ignore-lsp-errors
 capability flex-algo routing
 flex-algo 129
  metric-type link-delay
 exit-flex-algo
 !
 flex-algo 128
  metric-type te-metric
 exit-flex-algo
 !
 lsp-gen-interval 1
 max-lsp-lifetime 65535
 spf-interval-exp 50 5000
 metric-style wide
```

```
 mpls traffic-eng router-id 25.0.1.3
 mpls traffic-eng level-2
 dynamic-hostname
 set-overload-bit on-startup wait-for-bgp
 bfd all-interfaces
 net 49.0000.0100.0000.1033.00
 isis segment-routing global block 20000 80000
 segment-routing mpls
!
router bgp 65010
 neighbor 25.0.0.5 remote-as 65010
 neighbor 25.0.0.5 update-source loopback1
 neighbor 25.0.1.5 remote-as 65010
 neighbor 25.0.1.5 update-source loopback2
 !
 address-family ipv4 unicast
 neighbor 25.0.0.5 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 25.0.1.5 activate
 exit-address-family
 !
 exit
```

R4

```
qos enable
!
hostname RTR3
admin-group red 1
extended-admin-group green 32
extended-admin-group blue 93
!
router-id 25.0.1.3
!
segment-routing
 global block 20000 80000
!
interface ge2
 load-interval 30
 ip address 11.0.35.3/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface loopback1
 ip address 25.0.0.3/32 secondary
 ip router isis 1
!
```

```
interface loopback2
 ip address 25.0.1.3/32 secondary
 prefix-sid index 3 no-php
 prefix-sid algorithm-num 128 index 1283
 prefix-sid algorithm-num 129 index 1293
 ip router isis 1
!
interface xe4
 load-interval 30
 ip address 11.0.36.3/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface xe6
 speed 1g
 load-interval 30
 ip address 11.0.23.3/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
router isis 1
 is-type level-2-only
 ignore-lsp-errors
 capability flex-algo routing
 flex-algo 129
  metric-type link-delay
 exit-flex-algo
 !
 flex-algo 128
  metric-type te-metric
 exit-flex-algo
 !
 lsp-gen-interval 1
 max-lsp-lifetime 65535
 spf-interval-exp 50 5000
 metric-style wide
 mpls traffic-eng router-id 25.0.1.3
 mpls traffic-eng level-2
 dynamic-hostname
```

```
 set-overload-bit on-startup wait-for-bgp
 bfd all-interfaces
 net 49.0000.0100.0000.1033.00
 isis segment-routing global block 20000 80000
 segment-routing mpls
!
router bgp 65010
 neighbor 25.0.0.5 remote-as 65010
 neighbor 25.0.0.5 update-source loopback1
 neighbor 25.0.1.5 remote-as 65010
 neighbor 25.0.1.5 update-source loopback2
 !
 address-family ipv4 unicast
 neighbor 25.0.0.5 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 25.0.1.5 activate
 exit-address-family
 !
 exit
```

R5

```
qos enable
!
hostname RTR5
admin-group red 1
extended-admin-group green 32
extended-admin-group blue 93
router-id 25.0.1.5
!
segment-routing
 global block 20000 80000
!
interface ge2
 load-interval 30
 ip address 11.0.35.5/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface ge3
 load-interval 30
 ip address 11.0.25.5/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
```

```
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
!
interface loopback1
 ip address 25.0.0.5/32 secondary
 ip router isis 1
!
interface loopback2
 ip address 25.0.1.5/32 secondary
 prefix-sid index 5 no-php
 prefix-sid algorithm-num 128 index 1285
 prefix-sid algorithm-num 129 index 1295
 ip router isis 1
!
interface xe6
 speed 1g
 load-interval 30
 ip address 11.0.15.5/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface xe10
 load-interval 30
 ip address 11.0.56.5/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface xe17
 load-interval 30
 ip address 11.0.45.5/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
```

```
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
router isis 1
 is-type level-2-only
 ignore-lsp-errors
 capability flex-algo routing
 flex-algo 128
  metric-type te-metric
  priority 130
 exit-flex-algo
 !
 flex-algo 129
  metric-type link-delay
 exit-flex-algo
 !
 lsp-gen-interval 1
 max-lsp-lifetime 65535
 spf-interval-exp 50 5000
 metric-style wide
 mpls traffic-eng router-id 25.0.1.5
 mpls traffic-eng level-2
 dynamic-hostname
 set-overload-bit on-startup wait-for-bgp
 bfd all-interfaces
 net 49.0000.0100.0000.1055.00
 redistribute isis level-2 into level-1
 isis segment-routing global block 20000 80000
 segment-routing mpls
!
router bgp 65010
 neighbor 25.0.0.1 remote-as 65010
 neighbor 25.0.0.1 update-source loopback1
 neighbor 25.0.0.3 remote-as 65010
 neighbor 25.0.0.3 update-source loopback1
 neighbor 25.0.1.1 remote-as 65010
 neighbor 25.0.1.1 update-source loopback2
 neighbor 25.0.1.3 remote-as 65010
 neighbor 25.0.1.3 update-source loopback2
 !
 address-family ipv4 unicast
 neighbor 25.0.0.1 activate
 neighbor 25.0.0.3 activate
 exit-address-family
 !
 address-family vpnv4 unicast
 neighbor 25.0.1.1 activate
 neighbor 25.0.1.3 activate
 exit-address-family
 !
```

R6

```
qos enable
!
hostname RTR6
admin-group red 1
```

```
extended-admin-group green 32
extended-admin-group blue 93
router-id 25.0.1.6
!
segment-routing
 global block 20000 80000
!
interface loopback1
 ip address 25.0.0.6/32 secondary
 ip router isis 1
!
interface loopback2
 ip address 25.0.1.6/32 secondary
 prefix-sid index 6 no-php
 prefix-sid algorithm-num 128 index 1286
 prefix-sid algorithm-num 129 index 1296
 ip router isis 1
!
interface xe2
 load-interval 30
 ip address 11.0.36.6/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface xe10
 load-interval 30
 ip address 11.0.56.6/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo green
 isis extended-admin-group flex-algo blue
!
interface xe12
 load-interval 30
 ip address 11.0.26.6/24
 mtu 9216
 label-switching
 isis network point-to-point
 ip router isis 1
 isis te-metric flex-algo ipv4 10
 isis te-minimum-delay flex-algo 10
```

```
 isis te-maximum-delay flex-algo 20
 isis admin-group flex-algo red
 isis admin-group flex-algo anomaly red
 isis extended-admin-group flex-algo blue
!
router isis 1
 is-type level-2-only
 ignore-lsp-errors
 capability flex-algo routing
 flex-algo 128
  metric-type te-metric
 exit-flex-algo
 !
 flex-algo 129
  metric-type link-delay
 exit-flex-algo
 !
 lsp-gen-interval 1
 max-lsp-lifetime 65535
 spf-interval-exp 50 5000
 metric-style wide
 mpls traffic-eng router-id 25.0.1.6
 mpls traffic-eng level-2
 dynamic-hostname
 set-overload-bit on-startup wait-for-bgp
 bfd all-interfaces
 net 49.0000.0100.0000.1066.00
 isis segment-routing global block 20000 80000
 segment-routing mpls
```

TWAMP Configurations on R1:

```
hardware-profile filter twamp-ipv4 enable
!
delay-profile interfaces
 mode two-way
 burst-interval 1000
 burst-count 1
 interval 30
 advertisement periodic threshold 10
 advertisement periodic minimum-change 1000
 advertisement accelerated
 advertisement accelerated threshold 20
 advertisement accelerated minimum-change 2000
!
twamp-light reflector
 reflector-admin-state enable
 reflector-name RTR1-RTR2 reflector-ip ipv4 11.0.12.1 reflector-port 1025
!
twamp-light control
 control-admin-state enable
!
interface ge2
 delay-measurement dynamic twamp reflector-ip 11.0.12.2 reflector-port 1026
sender-ip 11.0.12.1
 loss-measurement dynamic
```

TWAMP Configurations on R2:

```
hardware-profile filter twamp-ipv4 enable
!
delay-profile interfaces
 mode two-way
 burst-interval 1000
 burst-count 1
 interval 30
 advertisement periodic threshold 10
 advertisement periodic minimum-change 1000
 advertisement accelerated
 advertisement accelerated threshold 20
 advertisement accelerated minimum-change 2000
!
twamp-light reflector
 reflector-admin-state enable
 reflector-name RTR1-RTR2 reflector-ip ipv4 11.0.12.2 reflector-port 1026
!
twamp-light control
 control-admin-state enable
!
interface ge2
 delay-measurement dynamic twamp reflector-ip 11.0.12.1 reflector-port 1025
sender-ip 11.0.12.2
 loss-measurement dynamic
```

Validation

Verify that IS-IS adjacencies are established with the expected neighbors and are operating within the Flex-Algo topology.

```
RTR1#show clns neighbors

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 3
Total number of adjacencies: 3
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
RTR2           ge2         e8c5.7ad4.7205      Up     24        L2   IS-IS
RTR5           ge6         e8c5.7a90.e1c8      Up     24        L2   IS-IS
RTR4           xe12        5c07.5828.af60      Up     24        L2   IS-IS
RTR1#sh isis topology

Tag 1:  VRF : default
IS-IS paths to level-2 routers
System Id            Metric     Next-Hop             Interface   SNPA
RTR1                 --
RTR2                 10         RTR2       ge2         e8c5.7ad4.7205
RTR3                 20         RTR2       ge2         e8c5.7ad4.7205
                                RTR5       ge6         e8c5.7a90.e1c8
RTR4                 10         RTR4       xe12        5c07.5828.af60
RTR5                 10         RTR5       ge6         e8c5.7a90.e1c8
RTR6                 20         RTR2       ge2         e8c5.7ad4.7205
```

```
                                RTR5       ge6         e8c5.7a90.e1c8
RTR1#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 1:  VRF : default
Total number of routes: 23

     Destination         Metric      Next-Hop            Interface      Tag
C    11.0.12.0/24        10          --                  ge2            0
C    11.0.14.0/24        10          --                  xe12           0
C    11.0.15.0/24        10          --                  ge6            0
L2   11.0.23.0/24        20          11.0.12.2           ge2            0
L2   11.0.24.0/24        20          11.0.14.4           xe12           0
                                      11.0.12.2           ge2            0
L2   11.0.25.0/24        20          11.0.12.2           ge2            0
                                      11.0.15.5           ge6            0
L2   11.0.26.0/24        20          11.0.12.2           ge2            0
L2   11.0.35.0/24        20          11.0.15.5           ge6            0
L2   11.0.36.0/24        30          11.0.12.2           ge2            0
                                      11.0.15.5           ge6            0
L2   11.0.45.0/24        20          11.0.14.4           xe12           0
                                      11.0.15.5           ge6            0
L2   11.0.56.0/24        20          11.0.15.5           ge6            0
C    25.0.0.1/32         10          --                  loopback1      0
L2   25.0.0.2/32         20          11.0.12.2           ge2            0
L2   25.0.0.3/32         30          11.0.12.2           ge2            0
                                      11.0.15.5           ge6            0
L2   25.0.0.4/32         20          11.0.14.4           xe12           0
L2   25.0.0.5/32         20          11.0.15.5           ge6            0
L2   25.0.0.6/32         30          11.0.12.2           ge2            0
                                      11.0.15.5           ge6            0
C    25.0.1.1/32         10          --                  loopback2      0
L2   25.0.1.2/32         20          11.0.12.2           ge2            0
L2   25.0.1.3/32         30          11.0.12.2           ge2            0
                                      11.0.15.5           ge6            0
L2   25.0.1.4/32         20          11.0.14.4           xe12           0
L2   25.0.1.5/32         20          11.0.15.5           ge6            0
L2   25.0.1.6/32         30          11.0.12.2           ge2            0
                                      11.0.15.5           ge6            0
```

Verify that the route for `25.0.1.3/32` is installed with the expected next-hop and metric, adhering to the Flex-Algo path constraints.

```
RTR1#sh ip route 25.0.1.3/32
VRF: Default, Routing entry for 25.0.1.3/32
  Known via "isis", distance 115, metric 30,  External Route Tag: 0, installed
00:38:27, best
  Last update 00:38:27 ago
```

```
  * 11.0.15.5, via ge6
  * 11.0.12.2, via ge2

RTR1#
RTR1#show ip isis route prefix 25.0.1.3/32 detail

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 1:  VRF : default
     Destination         Metric      ILM-ID     FTN-ID     In-Label   Next-Hop
Interface      Tag       Out-Label
L2   25.0.1.3/32         30          15         9          20003      11.0.12.2
ge2            0         20003
                                                          11.0.15.5         ge6
0         20003
     Src: 0100.0000.1033 Ifindex 10015
     Src: 0100.0000.1033 Ifindex 10003
```

Verify that IS-IS has SR enabled and supports the expected Flex-Algo capabilities, including SR-MPLS and algorithm-specific constraints.

```
RTR1#show isis segment-routing capability

Tag 1  Segment-Routing:
-----------------------------------------------------
Advertisement Router Capability   :25.0.1.1
Algorithm0                        :0
Algorithm1                        :129
Algorithm2                        :128
SRMS Preference                   :0
Total SID'S Supported             :60001
SR ERLD                           :6
SID Range List Count              :1
SID's Range                       :20000 - 80000
Total SID's Supported (SRLB)      :0
SRLB Range List Count             :0
-----------------------------------------------------
Advertisement Router Capability   :25.0.1.2
Algorithm0                        :0
Algorithm1                        :128
Algorithm2                        :129
SRMS Preference                   :0
Total SID'S Supported             :60001
SR ERLD                           :6
SID Range List Count              :1
SID's Range                       :20000 - 80000
Total SID's Supported (SRLB)      :0
SRLB Range List Count             :0
-----------------------------------------------------
Advertisement Router Capability   :25.0.1.3
```

```
Algorithm0                        :0
Algorithm1                        :129
Algorithm2                        :128
SRMS Preference                   :0
Total SID'S Supported             :60001
SR ERLD                           :6
SID Range List Count              :1
SID's Range                       :20000 - 80000
Total SID's Supported (SRLB)      :0
SRLB Range List Count             :0
-----------------------------------------------------
Advertisement Router Capability   :25.0.1.4
Algorithm0                        :0
Algorithm1                        :128
Algorithm2                        :129
SRMS Preference                   :0
Total SID'S Supported             :60001
SR ERLD                           :10
SID Range List Count              :1
SID's Range                       :20000 - 80000
Total SID's Supported (SRLB)      :0
SRLB Range List Count             :0
-----------------------------------------------------
Advertisement Router Capability   :25.0.1.5
Algorithm0                        :0
Algorithm1                        :128
Algorithm2                        :129
SRMS Preference                   :0
Total SID'S Supported             :60001
SR ERLD                           :6
SID Range List Count              :1
SID's Range                       :20000 - 80000
Total SID's Supported (SRLB)      :0
SRLB Range List Count             :0
-----------------------------------------------------
Advertisement Router Capability   :25.0.1.6
Algorithm0                        :0
Algorithm1                        :128
Algorithm2                        :129
SRMS Preference                   :0
Total SID'S Supported             :60001
SR ERLD                           :10
SID Range List Count              :1
SID's Range                       :20000 - 80000
Total SID's Supported (SRLB)      :0
SRLB Range List Count             :0
-----------------------------------------------------

RTR1#show isis segment-routing mapping-table ipv4 active
Tag 1  Segment-Routing:
Conflict Resolution Policy: Quarantine
```

```
Prefix              Range       Flags          Algo/SID-Index/Prefix-Flag List
25.0.1.1/32         1                            Algo:0   SID:1     PF:60
                                                 Algo:128 SID:1281  PF:40
                                                 Algo:129 SID:1291  PF:40
25.0.1.2/32         1                            Algo:0   SID:2     PF:60
                                                 Algo:128 SID:1282  PF:40
                                                 Algo:129 SID:1292  PF:40
25.0.1.3/32         1                            Algo:0   SID:3     PF:60
                                                 Algo:128 SID:1283  PF:40
                                                 Algo:129 SID:1293  PF:40
25.0.1.4/32         1                            Algo:0   SID:4     PF:60
                                                 Algo:128 SID:1284  PF:40
                                                 Algo:129 SID:1294  PF:40
25.0.1.5/32         1                            Algo:0   SID:5     PF:60
                                                 Algo:128 SID:1285  PF:40
                                                 Algo:129 SID:1295  PF:40
25.0.1.6/32         1                            Algo:0   SID:6     PF:60
                                                 Algo:128 SID:1286  PF:40
                                                 Algo:129 SID:1296  PF:40

MPLS Validation on RTR1
-----------------------------------
RTR1#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               FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-Intf
ELC       Nexthop         Algo-Num   UpTime
   i>   25.0.1.2/32       1        8        -          -    -          -         -
-               128        00:40:11
                                    1         0          Yes   3            ge2
No        11.0.12.2        -          -
   i>   25.0.1.2/32       3        8        -          -    -          -         -
-               129        00:40:11
                                    1         0          Yes   3            ge2
No        11.0.12.2        -          -
   i>   25.0.1.2/32       5        14       -          -    -          -         -
-               0          00:40:11
                                    13        0          Yes   20002        ge2
No        11.0.12.2        -          -
   i>   25.0.1.3/32       7        20       -          -    -          -         -
-               128        00:40:11
                                    29        0          Yes   21283        ge2
No        11.0.12.2        -          -
                                    19        0          Yes   21283        ge6
No        11.0.15.5        -          -
   i>   25.0.1.3/32       11       38       -          -    -          -         -
-               129        00:40:11
```

```
                                        37        0          Yes   21293       ge2
No        11.0.12.2       -          -
   i>   25.0.1.3/32        9        27       -         -    -          -         -
-               0          00:40:11
                                        32        0          Yes   20003       ge2
No        11.0.12.2       -          -
                                        26        0          Yes   20003       ge6
No        11.0.15.5       -          -
   i>   25.0.1.4/32        2        11       -         -    -          -         -
-               128        00:40:11
                                        22        0          Yes   3           xe12
No        11.0.14.4       -          -
   i>   25.0.1.4/32        4        11       -         -    -          -         -
-               129        00:40:11
                                        22        0          Yes   3           xe12
No        11.0.14.4       -          -
   i>   25.0.1.4/32        6        17       -         -    -          -         -
-               0          00:40:11
                                        16        0          Yes   20004       xe12
No        11.0.14.4       -          -
   i>   25.0.1.5/32        8        24       -         -    -          -         -
-               128        00:40:11
                                        9         0          Yes   3           ge6
No        11.0.15.5       -          -
   i>   25.0.1.5/32        12       24       -         -    -          -         -
-               129        00:40:11
                                        9         0          Yes   3           ge6
No        11.0.15.5       -          -
   i>   25.0.1.5/32        10       33       -         -    -          -         -
-               0          00:40:11
                                        31        0          Yes   20005       ge6
No        11.0.15.5       -          -
   i>   25.0.1.6/32        13       41       -         -    -          -         -
-               129        00:40:11
                                        40        0          Yes   21296       ge2
No        11.0.12.2       -          -
   i>   25.0.1.6/32        14       44       -         -    -          -         -
-               128        00:40:11
                                        43        0          Yes   21286       ge6
No        11.0.15.5       -          -
   i>   25.0.1.6/32        15       47       -         -    -          -         -
-               0          00:40:11
                                        35        0          Yes   20006       ge2
No        11.0.12.2       -          -
                                        46        0          Yes   20006       ge6
No        11.0.15.5       -          -
RTR1#show mpls forwarding-table 25.0.1.3/32
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
```

```
   i>   25.0.1.3/32        7        20       -          -    -          -         -
-                128         00:40:42
                                       29        0           Yes   21283        ge2
No         11.0.12.2        -           -
                                       19        0           Yes   21283        ge6
No         11.0.15.5        -           -
   i>   25.0.1.3/32        11       38       -          -    -          -         -
-                129         00:40:42
                                       37        0           Yes   21293        ge2
No         11.0.12.2        -           -
   i>   25.0.1.3/32        9        27       -          -    -          -         -
-                0           00:40:42
                                       32        0           Yes   20003        ge2
No         11.0.12.2        -           -
                                       26        0           Yes   20003        ge6
No         11.0.15.5        -           -
```

Verify that the FTN entry correctly maps the prefix 25.0.1.3/32 to the expected SR Label and next-hop based on the Flex-Algo policy.

```
RTR1#show mpls ftn-table 25.0.1.3/32
 Primary FTN entry with FEC: 25.0.1.3/32, id: 7, row status: Active, Tunnel-Policy: N/A,
State: Installed
  CreateTime: 00:40:48, UpTime: 00:40:48, LastUpdate: N/A
  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: 10, in intf: - in label: 0 out-segment ix: 29 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 29, owner: ISIS-SR-FA, Stale: NO, refcount: 4, out intf:
ge2, out label: 21283
    Nexthop addr: 11.0.12.2        cross connect ix: 10, op code: Push


     Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 19 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 19, owner: ISIS-SR-FA, Stale: NO, refcount: 2, out intf:
ge6, out label: 21283
    Nexthop addr: 11.0.15.5        cross connect ix: 10, op code: Push



 Primary FTN entry with FEC: 25.0.1.3/32, id: 11, row status: Active, Tunnel-Policy: N/
A, State: Installed
  CreateTime: 00:40:48, UpTime: 00:40:48, LastUpdate: N/A
  Owner: ISIS-SR-FA, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0,
Incoming DSCP: none, Algorithm Number:129
  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: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 37, owner: ISIS-SR-FA, Stale: NO, refcount: 2, out intf:
ge2, out label: 21293
    Nexthop addr: 11.0.12.2        cross connect ix: 9, op code: Push
```

```
 Primary FTN entry with FEC: 25.0.1.3/32, id: 9, row status: Active, Tunnel-Policy: N/A,
State: Installed
  CreateTime: 00:40:48, UpTime: 00:40:48, LastUpdate: N/A
  Owner: ISIS-SR, distance: 115, 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: 32 refcount: 1
      Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 32, owner: ISIS-SR, Stale: NO, refcount: 4, out intf: ge2,
out label: 20003
    Nexthop addr: 11.0.12.2        cross connect ix: 11, op code: Push

     Cross connect ix: 11, in intf: - in label: 0 out-segment ix: 26 refcount: 1
      Owner: ISIS-SR, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 26, owner: ISIS-SR, Stale: NO, refcount: 2, out intf: ge6,
out label: 20003
    Nexthop addr: 11.0.15.5        cross connect ix: 11, op code: Push
```

Verify the IS-IS path metrics assigned to each link in the Flex-Algo topology.

```
 RTR1#show isis topology algorithm 128

Tag 1:  VRF : default
IS-IS paths to level-2 routers
Flex-algo 128
System Id            Metric     Next-Hop             Interface   SNPA
RTR1                 --
RTR2                 10         RTR2       ge2         e8c5.7ad4.7205
RTR3                 20         RTR2       ge2         e8c5.7ad4.7205
                                RTR5       ge6         e8c5.7a90.e1c8
RTR4                 10         RTR4       xe12        5c07.5828.af60
RTR5                 10         RTR5       ge6         e8c5.7a90.e1c8
RTR6                 20         RTR5       ge6         e8c5.7a90.e1c8
```

Verify the flex algorithm 128 to view its configured parameters.

```
RTR1#show ip isis route algorithm 128

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 1:  VRF : default
Total number of routes: 23

     Destination         Metric      Next-Hop            Interface      Tag
C    11.0.12.0/24        10          --                  ge2            0
C    11.0.14.0/24        10          --                  xe12           0
C    11.0.15.0/24        10          --                  ge6            0
```

```
L2   11.0.23.0/24        20          11.0.12.2           ge2             0
L2   11.0.24.0/24        20          11.0.14.4           xe12            0
                                     11.0.12.2           ge2             0
L2   11.0.25.0/24        20          11.0.12.2           ge2             0
                                     11.0.15.5           ge6             0
L2   11.0.26.0/24        20          11.0.12.2           ge2             0
L2   11.0.35.0/24        20          11.0.15.5           ge6             0
L2   11.0.36.0/24        30          11.0.12.2           ge2             0
                                     11.0.15.5           ge6             0
L2   11.0.45.0/24        20          11.0.14.4           xe12            0
                                     11.0.15.5           ge6             0
L2   11.0.56.0/24        20          11.0.15.5           ge6             0
C    25.0.0.1/32         10          --                  loopback1       0
L2   25.0.0.2/32         20          11.0.12.2           ge2             0
L2   25.0.0.3/32         30          11.0.12.2           ge2             0
                                     11.0.15.5           ge6             0
L2   25.0.0.4/32         20          11.0.14.4           xe12            0
L2   25.0.0.5/32         20          11.0.15.5           ge6             0
L2   25.0.0.6/32         30          11.0.15.5           ge6             0
C    25.0.1.1/32         10          --                  loopback2       0
L2   25.0.1.2/32         20          11.0.12.2           ge2             0
L2   25.0.1.3/32         30          11.0.12.2           ge2             0
                                     11.0.15.5           ge6             0
L2   25.0.1.4/32         20          11.0.14.4           xe12            0
L2   25.0.1.5/32         20          11.0.15.5           ge6             0
L2   25.0.1.6/32         30          11.0.15.5           ge6             0
RTR1#
RTR1#show ip isis route prefix 25.0.1.3/32 algorithm 128

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 1:  VRF : default
     Destination         Metric      Next-Hop            Interface       Tag
L2   25.0.1.3/32         30          11.0.12.2           ge2             0
                                     11.0.15.5           ge6             0
RTR1#
```

Verify that the configured Flex-Algo parameters, constraints, and metrics are correctly applied and operational within the IS-IS domain.

```
RTR1#show isis flex-algo control
ISIS Instance : 1
 Num of times NSM disconnects received: 0
 Num of times NSM Admin Group mapping received: 2
 Num of times NSM Extended Admin Group mapping received: 3
 Flex-Algo ASLA advertising: On
 Flex-Algo strict ASLA usage: Off
 Flex-Algo Routing Capability: Yes
 Flex-Algo Readvertise Capability: No
 Flex-Algo CSPF Capability: Yes
```

```
 Flex-Algo Adjacency-SID Capability: No
 Flex-Algo Strict Adjacency-SID Capability: No
 Flex-Algo Backup Adjacency-SID Capability: No
 Flex-Algo QoS-Policy Capability: No
 Flex-Algo BGP-LS Capability: No
RTR1#
RTR1#show isis flex-algo 128 status winner detail
ISIS Instance : 1
  Router-ID : 25.0.1.2, System-ID : 0100.0000.1022
  Algorithm Number   : 128
  Metric Type        : te-metric
  Calculation Type   : spf
  Priority           : 131 (Winner)
  Prefix Metric Flag : No
  Fwd-Exclude-Any-AG Mask  : None
  Fwd-Include-Any-AG Mask  : None
  Fwd-Include-All-AG Mask  : None
  Rev-Exclude-Any-AG Mask  : None
  Rev-Include-Any-AG Mask  : None
  Rev-Include-All-AG Mask  : None
  Fwd-Exclude-Any-EAG Mask  : NULL
  Fwd-Include-Any-EAG Mask  :
    (32)
  Fwd-Include-All-EAG Mask  : NULL
  Rev-Exclude-Any-EAG Mask  : NULL
  Rev-Include-Any-EAG Mask  : NULL
  Rev-Include-All-EAG Mask  : NULL
  Exclude SRLG-ID List : NULL
  Exclude Minimum Bandwidth : 0
  Exclude Maximum Delay     : 0
  Intf-Group Mode Total B/W Calculate : No
  Reference Bandwidth : 0
  Granularity Bandwidth : 0

RTR1#
RTR1#
```

Verify that Flex-Algo 128 is enabled, correctly configured, and active with the expected user-defined constraints and parameters.

```
RTR1#show isis flex-algo 128 status usercfg summary
ISIS Instance : 1
  Algorithm Number  : 128
  Metric Type       : te-metric
  Calculation Type  : spf
  Priority          : 5
  Prefix Metric Flag: No
RTR1#
RTR1#show isis flex-algo 128 status election summary
ISIS Instance : 1
  Router-ID : 25.0.1.2, System-ID : 0100.0000.1022
  Algorithm Number   : 128
```

```
  Metric Type        : te-metric
  Calculation Type   : spf
  Priority           : 131 (Winner)
  Prefix Metric Flag : No

  Router-ID : 25.0.1.5, System-ID : 0100.0000.1055
  Algorithm Number   : 128
  Metric Type        : te-metric
  Calculation Type   : spf
  Priority           : 130
  Prefix Metric Flag : No

  Router-ID : 25.0.1.6, System-ID : 0100.0000.1066
  Algorithm Number   : 128
  Metric Type        : te-metric
  Calculation Type   : spf
  Priority           : 5
  Prefix Metric Flag : No

  Router-ID : 25.0.1.4, System-ID : 0100.0000.1044
  Algorithm Number   : 128
  Metric Type        : te-metric
  Calculation Type   : spf
  Priority           : 5
  Prefix Metric Flag : No

  Router-ID : 25.0.1.3, System-ID : 0100.0000.1033
  Algorithm Number   : 128
  Metric Type        : te-metric
  Calculation Type   : spf
  Priority           : 5
  Prefix Metric Flag : No

  Router-ID : 25.0.1.1, System-ID : 0100.0000.1011
  Algorithm Number   : 128
  Metric Type        : te-metric
  Calculation Type   : spf
  Priority           : 5
  Prefix Metric Flag : No
RTR1#show isis flex-algo 128 status winner summary
ISIS Instance : 1
  Router-ID : 25.0.1.2, System-ID : 0100.0000.1022
  Algorithm Number   : 128
  Metric Type        : te-metric
  Calculation Type   : spf
  Priority           : 131 (Winner)
  Prefix Metric Flag : No

RTR1#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: Disabled
Code    FEC               FTN-ID    Nhlfe-ID  Tunnel-ID   Pri   Out-Label    Out-Intf
ELC       Nexthop         Algo-Num   UpTime
   i>   25.0.1.2/32       1        8        -          -    -          -         -
-               128        01:19:19
                                    1         0          Yes   3            ge2
No        11.0.12.2        -          -
   i>   25.0.1.3/32       7        20       -          -    -          -         -
-               128        01:19:19
                                    29        0          Yes   21283        ge2
No        11.0.12.2        -          -
                                    19        0          Yes   21283        ge6
No        11.0.15.5        -          -
   i>   25.0.1.4/32       2        11       -          -    -          -         -
-               128        01:19:19
                                    22        0          Yes   3            xe12
No        11.0.14.4        -          -
   i>   25.0.1.5/32       8        24       -          -    -          -         -
-               128        01:19:19
                                    9         0          Yes   3            ge6
No        11.0.15.5        -          -
   i>   25.0.1.6/32       14       44       -          -    -          -         -
-               128        01:19:19
                                    43        0          Yes   21286        ge6
No        11.0.15.5        -          -
RTR1#
RTR1#show mpls ftn-table algorithm 128
 Primary FTN entry with FEC: 25.0.1.2/32, id: 1, row status: Active, Tunnel-Policy: N/A,
State: Installed
  CreateTime: 01:19:35, UpTime: 01:19:35, LastUpdate: N/A
  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: 2, in intf: - in label: 0 out-segment ix: 1 refcount: 1
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 1, owner: N/A, Stale: NO, refcount: 9, out intf: ge2, out
label: 3
    Nexthop addr: 11.0.12.2        cross connect ix: 2, op code: Push


 Primary FTN entry with FEC: 25.0.1.3/32, id: 7, row status: Active, Tunnel-Policy: N/A,
State: Installed
  CreateTime: 01:19:35, UpTime: 01:19:35, LastUpdate: N/A
  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: 10, in intf: - in label: 0 out-segment ix: 29 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
```

```
       Out-segment with ix: 29, owner: ISIS-SR-FA, Stale: NO, refcount: 4, out intf:
ge2, out label: 21283
    Nexthop addr: 11.0.12.2        cross connect ix: 10, op code: Push

     Cross connect ix: 10, in intf: - in label: 0 out-segment ix: 19 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 19, owner: ISIS-SR-FA, Stale: NO, refcount: 2, out intf:
ge6, out label: 21283
    Nexthop addr: 11.0.15.5        cross connect ix: 10, op code: Push

 Primary FTN entry with FEC: 25.0.1.4/32, id: 2, row status: Active, Tunnel-Policy: N/A,
State: Installed
  CreateTime: 01:19:35, UpTime: 01:19:35, LastUpdate: N/A
  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: 6, in intf: - in label: 0 out-segment ix: 22 refcount: 1
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 22, owner: N/A, Stale: NO, refcount: 9, out intf: xe12, out
label: 3
    Nexthop addr: 11.0.14.4        cross connect ix: 6, op code: Push

 Primary FTN entry with FEC: 25.0.1.5/32, id: 8, row status: Active, Tunnel-Policy: N/A,
State: Installed
  CreateTime: 01:19:35, UpTime: 01:19:35, LastUpdate: N/A
  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: 4, in intf: - in label: 0 out-segment ix: 9 refcount: 1
      Owner: N/A, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 9, owner: N/A, Stale: NO, refcount: 9, out intf: ge6, out
label: 3
    Nexthop addr: 11.0.15.5        cross connect ix: 4, op code: Push

 Primary FTN entry with FEC: 25.0.1.6/32, id: 14, row status: Active, Tunnel-Policy: N/
A, State: Installed
  CreateTime: 01:19:35, UpTime: 01:19:35, LastUpdate: N/A
  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: 13, in intf: - in label: 0 out-segment ix: 43 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 43, owner: ISIS-SR-FA, Stale: NO, refcount: 2, out intf:
ge6, out label: 21286
    Nexthop addr: 11.0.15.5        cross connect ix: 13, op code: Push
```

Implementation Examples

Traffic Optimization in Multi-Service Networks: Using flex algo bulk data transfers can prioritize cost-effective paths over low-latency ones.

Dynamic Resource Allocation in 5G Networks: In a 5G network, diverse applications such as real-time communication, IoT devices, and large data transfers require unique routing and resource allocation to meet their performance needs.Algo allows tailored routing paths for different applications (for example: low-latency routes for real-time communication, energy-efficient paths for IoT devices). Ensures optimal resource utilization while meeting the performance requirements of various 5G use cases.

Content Delivery Networks (CDNs): A CDN provider delivers content to end-users from multiple distributed servers. Flex Algo optimizes traffic flow to ensure the shortest delivery times based on user location and server availability.

CLI Commands

The `isis flex-algo` introduces the following configuration commands.

- `advertise-definition`
- `admin-group anomaly`
- `affinity-map`
- `affinity-ag-exclude-any`
- `affinity-ag-include-any`
- `affinity-ag-include-all`
- `affinity-ag-reverse-exclude-any`
- `affinity-ag-reverse-include-any`
- `affinity-ag-reverse-include-all`
- `affinity-eag-exclude-any`
- `affinity-eag-include-any`
- `affinity-eag-include-all`
- `affinity-eag-reverse-exclude-any`
- `affinity-eag-reverse-include-any`
- `affinity-ag-reverse-include-all`
- `asla flex-algo`
- `capability flex-algo`
- `exclude-maximum-delay`
- `extended-admin-group`
- `extended-admin-group anomaly`
- `flex-algo`
- `isis admin-group flex-algo`
- `isis admin-group anomaly flex-algo`
- `isis extended-admin-group flex-algo`
- `isis extended-admin-group anomaly flex-algo`

- isis te-metric flex-algo ipv4
- isis te-minimum-delay
- isis te-maximum-delay
- isis te-minimum-delay flex-algo
- isis te-maximum-delay flex-algo
- metric-type
- metric-type
- priority
- participate
- ti-lfa
- show isis flex-algo
- show isis flex-algo number status
- show isis flex-algo all status usercfg
- show isis flex-algo all status election
- show isis extend-admin-groups
- show isis flex-algo all status
- show isis flex-algo all status usercfg
- show isis flex-algo all status election
- show isis flex-algo all status winner
- show isis extend-admin-groups
- show isis extend-admin-groups

capability flex-algo

Use this command to enable the Flexible Algorithm feature on a per ISIS routing instance basis.

Use the `no` form of this command to disable the Flexible Algorithm feature in NSM

Command Syntax

```
capability flex-algo(routing)
(no) capability flex-algo (routing)
```

Parameters

`routing`	Enables the Flexible Algorithm feature on a per ISIS routing instance basis

Default

Disabled

Command Mode

Segment-routing Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
ocnos(config-router)#capability flex-algo routing
```

asla flex-algo

Use this command to enables the advertisement of Flexible algorithm specific ASLA link attributes in ISIS routing.

Use the `no` form of this command to disable the advertisement of Flexible algorithm specific ASLA link attributes in ISIS routing.

Command Syntax

```
asla flex-algo (advertise|strict)
(no) asla flex-algo (advertise|strict)
```

Parameters

`advertise`	Enables the advertisement of Flexible algorithm specific ASLA link attributes in ISIS routingx
`strict`	Enables the usage of strict ASLA link attributes for Flexible Algorithm specific path calculation in ISIS

Default

Disabled

Command Mode

Segment-routing Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
ocnos(config-router)#capability flex-algo routing
```

flex-algo

Use this command to create a new sub-mode `config-isis-fa` under ISIS router configuration mode.

Use the `no` form of this command to disable the new sub-mode `config-isis-fa` under ISIS router configuration mode.

Command Syntax

```
flex-algo <Algorithm-Number>
```

Parameters

`Algorithm-Number`	Specifies the flexible algorithm number <128-255>.

Default

Disabled

Command Mode

ISIS Router Configuration Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
ocnos(config-router)#flex-algo 128
```

isis admin-group flex-algo

Use this command to assign an Admin Group to a specific Flex-Algorithm (Flex-Algo) in IS-IS .

Use the `no` form of this command to assign an Admin Group to a specific Flex-Algorithm (Flex-Algo) in IS-IS. .

Command Syntax

```
isis admin-group flex-algo <1 - 16777214>
```

Parameters

`1 - 16777214`	Specifies the Flex-Algorithm ID that will be associated with an Admin Group.

Default

Disabled

Command Mode

ISIS Router Configuration Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# admin-group flex-algo 100
```

affinity-map

Use this command to command associates the affinity name with the particular bit positions in the Extended Admin Group (EAG) bitmask.

The Affinity-Map (EAG) definition must be globally unique across all routers in a SR domain.

Use the `no` form of this command to affinity name with the particular bit positions.

Command Syntax

```
affinity-mapp (attribute|Bit-Number)
```

Parameters

`attribute`	Specifies the name of the affinity map.
`Bit-Number`	Bit position in the Extended Admin Group bitmask. Minimum value is 32 and maximum value is 95.

Default

Disabled

Command Mode

ISIS Router Configuration Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
ocnos(config-router)#affinity-mapping
```

metric-type

Use this command to specify the type of metric to be used for path computation in routing protocols, such as IS-IS or Traffic Engineering. This command allows selecting different metric types based on network optimization requirements.

Use the `no` form of this command to enable metric type as igp-metric.

Command Syntax

```
metric-type {igp-metric | te-metric | link-delay}
(no) metric-type {igp-metric | te-metric | link-delay}
```

Parameters

`igp-metric`	Uses the Interior Gateway Protocol (IGP) metric for routing decisions. .
`te-metric`	Uses the Traffic Engineering (TE) metric, which is specifically configured for MPLS TE.
`link-delay`	Uses the measured or configured link delay as the metric, allowing for delay-based path selection.

Default

IGP

Command Mode

ISIS Router Configuration Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
ocnos(config-isis-fa)#metric-type te-metric
```

priority

Use this command to define the priority value for a specific process, protocol, or routing decision. The priority value determines the precedence of an entity when multiple options exist.

Use the `no` form of this command is used then the priority value is updated to the default value of 5.

Command Syntax

```
priority <1-255>
(no) priority <1-255>
```

Parameters

`1-255`	Priority for Flexible Algorithm (nor for routing).

Default

5

Command Mode

ISIS Router Configuration Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
ocnos(config-isis-fa)# priority 200
```

advertise-definition

Use this command to enable advertisement of the Flexible Algorithm definition.

Use the `no` form of this command to disable advertisement of the Flexible Algorithm definition.

Command Syntax

```
advertise-definition
(no) advertise-definition
```

Parameters

None

Default

5

Command Mode

ISIS Router Configuration Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
ocnos(config-isis-fa)# advertise-definition
```

participate

Use this command to enable participation for a Flexible Algorithm.

Use the `no` form of this command to disable participation for a Flexible Algorithm.

Command Syntax

```
participate
(no) participate
```

Parameters

None

Default

None

Command Mode

ISIS Router Configuration Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
ocnos(config-isis-fa)# participate
```

ti-lfa

Use this command to enable TI-LFA based fast rerouting of primary paths associated to a flexible algorithm.

Use the `no` form of this command to disable TI-LFA based fast rerouting of primary paths associated to a flexible algorithm.

Command Syntax

```
ti-lfa
(no) ti-lfa
```

Parameters

None

Default

None

Command Mode

ISIS Router Configuration Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
ocnos(config-isis-fa)# ti-lfa
```

exclude-maximum-delay

Use this command to set the exclude maximum link delay value constraint for a Flexible Algorithm definition.

Use the `no` form of this command to unset the exclude maximum link delay value constraint for a Flexible Algorithm definition.

Command Syntax

```
exclude-maximum-delay <1-16777215>
(no)exclude-maximum-delay <1-16777215>
```

Parameters

`<1-16777215>`	The maximum delay value in the range.

Default

None

Command Mode

Interface Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-if)# exclude-maximum-delay 100000
```

affinity-ag-exclude-any

Use this command to configure a Exclude any Admin-Group affinity name for the forward direction.

Use the `no` form of this command to disable the exclude maximum link delay value constraint for a Flexible Algorithm definition.

Command Syntax

```
  affinity-ag-exclude-any <NAME>
(no)  affinity-ag-exclude-any <NAME>
```

Parameters

`<Name>`	Specifies the name of the affinity group to be excluded from path computation.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router mpls
OcNOS(config-mpls)# traffic-eng
OcNOS(config-mpls-te)# affinity-ag-exclude-any 0x02
```

affinity-ag-include-any

Use this command to configure a include any Admin-Group affinity name for the forward direction.

Use the `no` form of this command to exclude the to configure a Include any Admin-Group affinity name for the forward direction.

Command Syntax

```
affinity-ag-include-any <Affinity-Name>
(no) affinity-ag-include-any <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the affinity group(s) that must be included in the path.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te-po)# affinity-ag-include-any TRUSTED_PATHS
```

affinity-ag-include-all

Use this command to to configure a include all Admin-Group affinity name for the forward direction.

Use the `no` form of this command to exclude include all Admin-Group affinity name for the forward direction.

Command Syntax

```
 affinity-ag-include-all} <Affinity-Name>
(no) affinity-ag-include-all} <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the name of the Admin-Group affinity to be included.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te-po)# affinity-ag-include-all FORWARD_TRUSTED_PATHS
```

affinity-ag-reverse-exclude-any

Use this command to configure a exclude any Admin-Group affinity name for the reverse direction.

Use the `no` form of this command to disable a exclude any Admin-Group affinity name for the reverse direction.

Command Syntax

```
affinity-ag-reverse-exclude-any <Affinity-Name>
(no)  affinity-ag-reverse-exclude-any <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the name of the Admin-Group affinity to be excluded.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te)# affinity-ag-reverse-exclude-any REVERSE_HIGH_LATENCY
```

affinity-ag-reverse-include-any

Use this command to configure a include any Admin-Group affinity name for the reverse direction.

Use the `no` form of this command to exclude a include any Admin-Group affinity name for the reverse direction.

Command Syntax

```
affinity-ag-reverse-include-any <Affinity-Name>
(no) affinity-ag-reverse-include-any <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the name of the Admin-Group affinity to be included for the reverse direction.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te-po)# affinity-ag-reverse-include-any
REVERSE_TRUSTED_PATHS
```

affinity-ag-reverse-include-all

Use this command to configure a include all Admin-Group affinity name for the reverse direction.

Use the `no` form of this command to exclude a Include all Admin-Group affinity name for the reverse direction.

Command Syntax

```
   affinity-ag-reverse-include-all} <Affinity-Name>
(no) affinity-ag-reverse-include-all} <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the Admin-Group affinity name to include all associated links in the reverse direction.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te)# affinity-ag-reverse-include-all REVERSE_LOW_LATENCY
REVERSE_HIGH_PRIORITY
```

affinity-eag-exclude-any

Use this command to configure a exclude any Explicit Affinity Group (EAG) affinity name for the forward direction.

Use the `no` form of this command to disable a exclude any EAG affinity name for the forward direction.

Command Syntax

```
affinity-eag-exclude-any <Affinity-Name>
(no) affinity-eag-exclude-any <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the EAG affinity name that should be excluded from path selection.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te-po)# affinity-eag-exclude-any EAG_UNTRUSTED
```

affinity-eag-include-any

Use this command to configure a include any Explicit Affinity Group (EAG) affinity name for the forward direction.

Use the `no` form of this command to disable a include any EAG affinity name for the forward direction.

Command Syntax

```
   affinity-eag-include-any <Affinity-Name>
(no) affinity-eag-include-any <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the EAG affinity name that should be included from path selection.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te-po)# affinity-eag-include-any EAG_TRUSTED_PATHS
```

affinity-eag-include-all

Use this command to configure a include all Explicit Affinity Group (EAG) affinity name for the forward direction.

Use the `no` form of this command to disable a include all EAG affinity name for the forward direction.

Command Syntax

```
   affinity-eag-include-all <Affinity-Name>
(no) affinity-eag-include-all <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the EAG affinity name that should be included from path selection.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te)# affinity-eag-include-all EAG_LOW_LATENCY
EAG_HIGH_PRIORITY
```

affinity-eag-reverse-exclude-any

Use this command to configure a exclude any Explicit Affinity Group (EAG) affinity name for the reverse direction.

Use the `no` form of this command to disable a exclude any EAG affinity name for the reverse direction.

Command Syntax

```
affinity-eag-reverse-exclude-any <Affinity-Name>
(no) affinity-eag-reverse-exclude-any <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the EAG affinity name that should be excluded from reverse direction path computation.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te-po)# affinity-eag-reverse-exclude-any EAG_UNTRUSTED
```

affinity-eag-reverse-include-any

Use this command to configure a include any Explicit Affinity Group (EAG) affinity name for the reverse direction.

Use the `no` form of this command to disable a include any EAG affinity name for the reverse direction.

Command Syntax

```
affinity-eag-reverse-include-any <Affinity-Name>
(no) affinity-eag-reverse-include-any <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the EAG affinity name that should be included from reverse direction path computation.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te)# affinity-eag-reverse-include-any EAG_LOW_LATENCY
EAG_HIGH_PRIORITY
```

affinity-eag-reverse-include-all

Use this command to configure a include all Explicit Affinity Group (EAG) affinity name for the reverse direction.

Use the `no` form of this command to disable a include all EAG affinity name for the reverse direction.

Command Syntax

```
affinity-eag-reverse-include-all <Affinity-Name>
(no) affinity-eag-reverse-include-all <Affinity-Name>
```

Parameters

`<Affinity Name>`	Specifies the EAG affinity name that should be included all links from reverse direction path computation.

Default

None

Command Mode

MPLS Traffic Engineering (TE) configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config-mpls-te)# affinity-eag-reverse-include-all EAG_LOW_LATENCY
EAG_HIGH_PRIORITY
```

isis te-metric flex-algo ipv4

Use this command to set the Traffic Engineering (TE) metric on an interface for the Flexible Algorithm application..

Use the `no` form of this command to unset the TE metric on an interface for the Flexible Algorithm application..

Command Syntax

```
    isis te-metric flex-algo ipv4 <1-16777214>
(no)  isis te-metric flex-algo ipv4 <1-16777214>
```

Parameters

`<1-16777214>`	Specifies the Flex-Algorithm ID.

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# te-metric flex-algo ipv4 1
```

isis te-minimum-delay

Use this command to configure the Traffic Engineering (TE) maximum delay value on a ISIS routing enabled interface for the Flexible Algorithm application.

Use the `no` form of this command to disable the TE maximum delay value on a ISIS routing enabled interface for the Flexible Algorithm application.

Command Syntax

```
isis te-minimum-delay flex-algo <1-16777214>
(no) isis te-minimum-delay flex-algo <1-16777214>
```

Parameters

`<1-16777214>`	Specifies the Flex-Algorithm ID for which the minimum delay metric should be applied.

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# te-minimum-delay flex-algo 1
```

isis te-maximum-delay

Use this command to configure the Traffic Engineering (TE) maximum delay value on a ISIS routing enabled interface for the Flexible Algorithm application.

Use the `no` form of this command to disable the TE maximum delay value on a ISIS routing enabled interface for the Flexible Algorithm application.

Command Syntax

```
isis te-maximum-delay flex-algo <1-16777214>
(no) isis te-maximum-delay flex-algo <1-16777214>
```

Parameters

`<1-16777214>`	Specifies the Flex-Algorithm ID for which the maximum delay metric should be applied.

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# te-maximum-delay flex-algo 1
```

isis te-minimum-delay flex-algo

Use this command to configure the Traffic Engineering (TE) minimum delay value on a ISIS routing enabled interface for the Flexible Algorithm application.

Use the `no` form of this command to disable the TE minimum delay value on a ISIS routing enabled interface for the Flexible Algorithm application.

Command Syntax

```
        isis te-minimum-delay flex-algo <1-16777214>
(no)    isis te-minimum-delay flex-algo <1-16777214>
```

Parameters

`<1-16777214>`	Specifies the Flex-Algorithm ID that the minimum delay metric will be applied to.

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# te-minimum-delay flex-algo 1
```

isis te-maximum-delay flex-algo

Use this command to configure the Traffic Engineering (TE) maximum delay value on a ISIS routing enabled interface for the Flexible Algorithm application.

Use the `no` form of this command to disable the TE maximum delay value on a ISIS routing enabled interface for the Flexible Algorithm application.

Command Syntax

```
        isis te-maximum-delay flex-algo <1-16777214>
(no)    isis te-maximum-delay flex-algo <1-16777214>
```

Parameters

`<1-16777214>`	Specifies the Flex-Algorithm ID that the maximum delay metric will be applied to.

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# te-maximum-delay flex-algo 1
```

extended-admin-group

Use this command to associate an Extended Admin Group (EAG) name with a specific bit position in the Extended Admin Group bitmask.

Use the `no` form of this command to disassociate an EAG name with a specific bit position in the Extended Admin Group bitmask.

Command Syntax

```
extend-admin-group <Extended-Admin-Group-Name> bit-position <Bit-Number>
(no) extend-admin-group <Extended-Admin-Group-Name> bit-position <Bit-Number>
```

Parameters

`Affinity-Name`	Specifies the name of the affinity map.
`Bit-Number`	Specifies bit position in the Extended Admin Group bitmask. Minimum value is 32 and maximum value is 95.

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# extend-admin-group EAG_FAST_PATH bit-position 2
```

isis extended-admin-group flex-algo

Use this command to set the affinity, i.e. Extended Admin Group (EAG) on a routing enabled, interface for the Flexible Algorithm application.

Use the `no` form of this command to remove the affinity, i.e. EAG on a routing enabled, interface for the Flexible Algorithm application.

Command Syntax

```
isis extended-admin-group flex-algo <Affinity-Name>
(no) isis extended-admin-group flex-algo <Affinity-Name>
```

Parameters

`Affinity-Name`	Specifies the Flex-Algorithm ID that will be associated with an Extended Admin Group.

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# extended-admin-group flex-algo 10
```

extended-admin-group anomaly

Use this command to enable anomaly detection for the Extended Admin Group (EAG) in IS-IS.

Use the `no` form of this command to disable the enable anomaly detection for the Extended Admin Group (EAG) in IS-IS.

Command Syntax

```
        extended-admin-group anomaly
(no) extended-admin-group anomaly
```

Parameters

None

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# extended-admin-group anomaly
```

isis extended-admin-group anomaly flex-algo

Use this command to detect anomalies in the Extended Admin Group (EAG) configuration for a specific Flex-Algo in IS-IS.

Use the `no` form of this command to disable anomalies in the Extended Admin Group (EAG) configuration for a specific Flex-Algo in IS-IS.

Command Syntax

```
isis extended-admin-group anomaly flex-algo < 1 - 16777214>
(no) isis extended-admin-group anomaly flex-algo < 1 - 16777214>
```

Parameters

`<1 – 16777214>`	Specifies the Flex-Algorithm ID for which EAG anomaly detection should be enabled.

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# extended-admin-group anomaly flex-algo 100
```

admin-group anomaly

Use this command to enable anomaly detection for the Admin Group configuration in IS-IS.

Use the `no` form of this command to disable anomaly detection for the Admin Group configuration in IS-IS..

Command Syntax

```
admin-group anomaly
(no) admin-group anomaly
```

Parameters

None

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# admin-group anomaly
```

isis admin-group anomaly flex-algo

Use this command to detect anomalies in the Admin Group configuration for a specific Flex-Algo in IS-IS.

Use the `no` form of this command to disable anomalies in the Admin Group configuration for a specific Flex-Algo in IS-IS.

Command Syntax

```
isis admin-group anomaly flex-algo < 1 - 16777214>
(no) isis admin-group anomaly flex-algo < 1 - 16777214>
```

Parameters

`<1 – 16777214>`	Specifies the Flex-Algorithm ID to monitor for anomalies.

Default

None

Command Mode

ISIS Router configuration mode.

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS(config)# router isis
OcNOS(config-isis)# admin-group anomaly flex-algo 50
```

show isis flex-algo

Use this command to display information about the flexible algorithm configuration and operational state in IS-IS.

Command Syntax

```
show isis flex-algo (128-255>) (all) (control)
```

Parameters

`Algorithm-Number`	Specifies the flexible algorithm number <128-255>.
`all`	Specifies all flexible algorithms.
`control`	Specifies flexible algorithm global control.

Default

None

Command Mode

Execution Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

```
OcNOS#configure terminal
```

show isis flex-algo number status

Use this command to display status information for a specific Flexible Algorithm identified by `(128-255)`.

Command Syntax

```
show isis flex-algo (128-255) status
```

Parameters

`election`	Displays whether the algorithm is part of a group of algorithms being elected to handle the path computation.
`usercfg`	Displays the algorithms that have been manually configured by the user on the device..
`winner`	Displays which algorithm has been selected for path computation based on the election process.

Default

None

Command Mode

Execution Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

View Election Status for Flexible Algorithm:

```
show isis flex-algo 150 status election
```

View Election Status for Flexible Algorithm:

```
show isis flex-algo 150 status usercfg
```

View Election Status for Flexible Algorithm:

```
show isis flex-algo 150 status winner
```

show isis flex-algo all status usercfg

Use this command to display the details of flexible algorithm.

Command Syntax

```
show isis flex-algo all status usercfg (summary/detail))
```

Parameters

`summary`	Displays a summary of user configuration of all flexible algorithm(s) of local router node.
`detail`	Displays in detail user configuration of all flexible algorithm(s) of local router node.

Default

None

Command Mode

Execution Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

View isis Status for Flexible Algorithm:

```
show isis flex-algo 150 status usercfg summary
```

show isis flex-algo all status election

Use this command to sisplay a summary of all flexible algorithm's FAD learnt from all router(s), that participated in the election process.

Command Syntax

```
show isis flex-algo <128-255> all status election (summary/detail))
```

Parameters

`summary`	Displays a summary of all flexible algorithm's FAD learnt from all router(s), that participated in the election process.
`detail`	Displays in detail all flexible algorithm's FAD learnt from all router(s), that participated in the election process.

Default

None

Command Mode

Execution Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

View Election Status for Flexible Algorithm:

```
show isis flex-algo 128 all status election summary
```

```
show isis flex-algo 141 all status election detail
```

show isis flex-algo all status winner

Use this command to displays a summary all flexible algorithm's FAD that have been declared as election winner across all routers of the SR domain.

Command Syntax

```
show isis flex-algo <128-255> all status winner (summary/detail))
```

Parameters

`summary`	Displays a summary all flexible algorithm's FAD that have been declared as election winner across all routers of the SR domain.
`detail`	Displays in detail all flexible algorithm's FAD that have been declared as election winner across all routers of the SR domain.

Default

None

Command Mode

Execution Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

View Election Status for Flexible Algorithm:

```
show isis flex-algo 128 all status winner summary
```

```
show isis flex-algo 141 all status winner detail
```

show isis extend-admin-groups

Use this command to display the global extended admin group mapping table (Affinity map).

Command Syntax

```
show isis extend-admin-groups
```

Parameters

None

Default

None

Command Mode

Execution Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

View Election Status for Flexible Algorithm:

```
show isis extend-admin-groups
```

show isis flex-algo all status

Use this command to display the status of all Flexible Algorithms in the IS-IS protocol.

Command Syntax

```
show isis flex-algo all status (election/userconfig/winner(summary/detail))
```

Parameters

`election`	Displays whether the algorithm is part of a group of algorithms being elected to handle the path computation.
`usercfg`	Displays the algorithms that have been manually configured by the user on the device.
`winner`	Displays which algorithm has been selected for path computation based on the election process.
`summary`	Displays a summary of user configuration of flexible algorithm of local router node.
`detail`	Displays in detail user configuration of flexible algorithm of local router node.

Default

None

Command Mode

Execution Mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

View Election Status for all Flexible Algorithms:

```
show isis flex-algo 150 status election
```

View Election Status for all Flexible Algorithms:

```
show isis flex-algo 150 status usercfg
```

View Election Status for all Flexible Algorithms:

```
show isis flex-algo 150 status winner
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Intermediate System to Intermediate System (IS-IS)	A link-state routing protocol used to exchange routing information within a network. It is widely used in large enterprise and service provider networks.
Flexible Algorithm (Flex-Algo)	A feature in IS-IS that allows for the definition of custom path computation algorithms. It provides more granular control over traffic routing, enabling the use of application-specific routing decisions
Segment Routing (SR)	A type of source-based routing in which the sender specifies the route that a packet will take through the network by encoding the path as a sequence of segments. In IS-IS, SR can be enhanced with flexible algorithms to improve traffic engineering.
Link-State Advertisements(LSAs)	Packets exchanged between IS-IS routers that contain information about the network's topology, such as available links, node statuses, and routing information.

CHAPTER 4 Traffic Steering for Flexible Algorithms

Overview

Traffic Steering for Flexible Algorithms feature integrates BGP On-Demand Next Hop (ODN) policies with Flexible Algorithms to improve the efficiency of path computation. Flexible Algorithms allow for route selection based on specific network constraints such as latency, bandwidth, and other performance metrics. By combining BGP ODN policies with Flexible Algorithms, the network can dynamically compute and adjust optimal paths in real-time, ensuring effective traffic management according to service requirements.

Feature Characteristics

- **Coloring for Traffic Steering:**
  - The egress PE node assigns colors to MPLS service FTNs.
  - These colors represent SR-TE SLA requirements and are advertised via BGP UPDATE messages to the ingress PE node.
  - The ingress node matches the color information with the corresponding ODN policy and steers the traffic accordingly (Traffic Steering).
- **Integration with Flexible Algorithms:**
  - ODN policies can be combined with Flexible Algorithms to create Flex-Algo-based SR-MPLS BE tunnels (Flex-Algo LSPs).
  - These tunnels enhance path computation by allowing route selection based on constraints such as latency, bandwidth, and traffic load.
  - Network operators can customize routing logic to optimize traffic flow and enhance network efficiency.
- **Support for Multiple MPLS Services:**
  - Flex-Algo LSPs can be used with:
    - VPLS (Signaling via BGP)
    - BGP VPNv4
    - BGP VPNv6
    - EVPN (ELINE, ELAN, ETREE)

Benefits

This feature offers several key advantages:

- Traffic steering based on SR-TE SLA requirements, ensuring optimal path selection.
- Customizable routing logic using Flexible Algorithms, tailored to constraints like latency, bandwidth, and traffic load.
- Dynamic network adaptation, which reduces the need for manual intervention and minimizes complexity.
- Efficient resource utilization through Flex-Algo LSPs, optimizing traffic routing.
- Policy-driven traffic steering, enhancing resource allocation and overall network performance.
- Real-time path adjustments in response to congestion, failures, or changes in SLA.

Prerequisites

This feature requires the following device capabilities:

- OcNOS devices must support:
- ISIS-SR Flex-Algo and ECMP
- BGP-VPLS and L3VPN/6VPE services
- EVPN services

Configuration

The following configuration enables Flex Algo with ISIS-SR.

Topology

This topology represents the application of Flex-Algo in a SR network environment, with three distinct algorithms highlighted: Flex Algo 0, 128, and 129. Each algorithm defines a specific logical topology based on unique constraints and use cases.

Topology Visualization:

- Red Lines (Flex Algo 0): Represents the default IGP metric-based paths where all links contribute to general traffic forwarding.
- Green Lines (Flex Algo 128): Highlights paths optimized for bandwidth efficiency using GREEN-affinity links, filtering out non-compliant links.
- Blue Lines (Flex Algo 129): Displays paths optimized for delay-sensitive traffic using BLUE-affinity links, excluding all others.

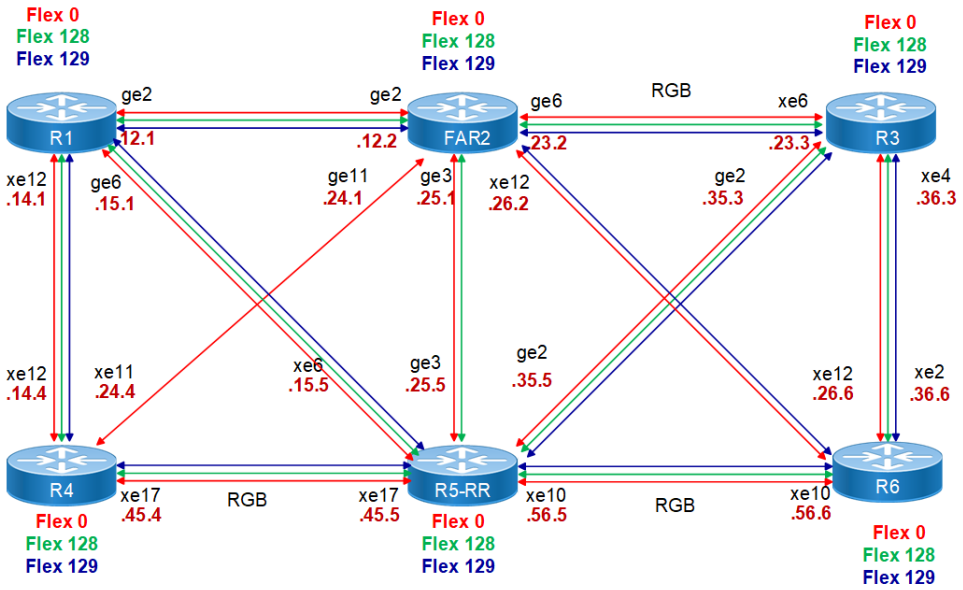

Logical Topologies Based on Flexible Algorithms

Configuration for BGP-VPLS Traffic Steering with ODN for FlexAlgo

The following configuration enables BGP-VPLS Traffic Steering with ODN Policy.

- The goal is to extend the existing setup used in the Flex Algorithm for ISIS chapter to configure BGP VPLS services over Flex-Algo 128.
- ODN policy will be used to automate the steering of VPLS traffic. This ensures that VPLS traffic follows the optimized paths based on Flex-Algo 128 constraints and metrics.
- The approach helps in seamless integration of VPLS over an SR-based network while maintaining traffic engineering flexibility.

To configure BGP-VPLS functionality on nodes between RTR1 to RTR3, follow the steps mentioned below:

1. Configure BGP for VPLS Services.

    1. Set up BGP in L2VPN address family on RTR1 and RTR2 as below.

```
RTR1(config)#router bgp 65010
RTR1(config-router)# bgp auto-policy-soft-reset enable
RTR1(config-router)# neighbor 25.0.1.5 remote-as 65010
RTR1(config-router)# neighbor 25.0.1.5 update-source loopback2
RTR1(config-router)# end

RTR1(config-router)# address-family l2vpn vpls
RTR1(config-router-af)# neighbor 25.0.1.5 activate
RTR1(config-router-af)# exit-address-family
RTR1(config-router)# end
RTR1(config-router)# exit

RTR3(config)#router bgp 65010
RTR3(config-router)# bgp auto-policy-soft-reset enable
RTR3(config-router)# neighbor 25.0.1.5 remote-as 65010
RTR3(config-router)# neighbor 25.0.1.5 update-source loopback2
RTR1(config-router)# end
RTR3(config-router)# exit

RTR3(config-router)# address-family l2vpn vpls
RTR3(config-router-af)# neighbor 25.0.1.5 activate
RTR3(config-router-af)# exit-address-family
RTR1(config-router)# end
RTR3(config-router)# exit
```

    2. Set up BGP in L2VPN address family on RTR3 which acts as RR below:

```
RTR5(config)#router bgp 65010
RTR5(config-router)# bgp auto-policy-soft-reset enable
RTR5(config-router)# no bgp inbound-route-filter
RTR5(config-router)# neighbor 25.0.1.1 remote-as 65010
RTR5(config-router)# neighbor 25.0.1.1 update-source loopback2
RTR5(config-router)# neighbor 25.0.1.3 remote-as 65010
RTR5(config-router)# neighbor 25.0.1.3 update-source loopback2
RTR5(config-router)# !
RTR5(config-router)# address-family l2vpn vpls
RTR5(config-router-af)# neighbor 25.0.1.1 activate
RTR5(config-router-af)# neighbor 25.0.1.1 route-reflector-client
RTR5(config-router-af)# neighbor 25.0.1.3 activate
```

```
RTR5(config-router-af)# neighbor 25.0.1.3 route-reflector-client
RTR5(config-router-af)# exit-address-family
RTR5(config-router)# !
RTR5(config-router)# exit
```

2. Configure VPLS Instance on Between RTR1 and RTR3.

```
RTR1(config)#mpls vpls BGP-VPLS 1000
RTR1(config-vpls)# signaling bgp
RTR1(config-vpls-sig)#  ve-id 1
RTR1(config-vpls-sig)#  exit-signaling
RTR1(config-vpls)# exit-vpls
RTR1(config)#end

RTR3(config)#mpls vpls BGP-VPLS 1000
RTR3(config-vpls)# signaling bgp
RTR3(config-vpls-sig)#  ve-id 1
RTR3(config-vpls-sig)#  exit-signaling
RTR3(config-vpls)# exit-vpls
RTR3(config)#end
```

Note: Each **VE-ID** must be unique per node within the VPLS instance.

3. Configure the egress side to advertise the color per VPLS instance using a route-map:

    1. Egress-side, Color needs to be advertised per vpls instance (as mentioned below).

```
RTR3(config)#route-map set_color permit 10
RTR3(config-route-map)# set extcommunity color 1000
RTR3(config-route-map)#!

RTR3(config)#
RTR3(config)#mpls vpls BGP-VPLS 1000
RTR3(config-vpls)#route-map set_color
RTR3(config-vpls)#exit
RTR3(config)#
```

4. Enable ODN Policy for Traffic Steering on RTR1:

```
RTR1(config)#segment-routing
RTR1(config-sr)# traffic-engineering
RTR1(config-sr-te)# on-demand-nexthop 1000
RTR1(config-sr-odn)# flex-algo 128
RTR1(config-sr-odn)# exit-sr-odn
RTR1(config-sr-te)# end
RTR1(config-sr-te)# exit-te
RTR1(config-sr)#
```

5. Configure ACCESS Interface for BGP-VPLS.

```
RTR1(config)#interface xe15
RTR1(config-if)# mtu 9216
RTR1(config)#interface xe15.1000 switchport
RTR1(config-if)# encapsulation dot1q 1000
RTR1(config-if)# mtu 9216
RTR1(config-if)# access-if-vpls
RTR1(config-acc-if-vpls)#  mpls-vpls BGP-VPLS
RTR1(config-acc-if-vpls)#!
```

```
RTR1(config)#interface xe15
RTR1(config-if)# mtu 9216
RTR1(config)#interface xe19.1000 switchport
RTR1(config-if)# encapsulation dot1q 1000
RTR1(config-if)# mtu 9216
RTR1(config-if)# access-if-vpls
RTR1(config-acc-if-vpls)#  mpls-vpls BGP-VPLS
RTR1(config-acc-if-vpls)#!
```

Configuration Sanapshot:

R1

```
mpls vpls BGP-VPLS 1000
signaling bgp
ve-id 1
exit-signaling
exit-vpls

interface xe15
 mtu 9216

interface xe15.1000 switchport
 encapsulation dot1q 1000
 mtu 9216
 access-if-vpls
 mpls-vpls BGP-VPLS

router bgp 65010
bgp auto-policy-soft-reset enable
address-family l2vpn vpls
 neighbor 25.0.1.5 activate
 exit-address-family

segment-routing
 traffic-engineering
  on-demand-nexthop 1000
   flex-algo 128
  exit-sr-odn
  !
 exit-te
```

R2

```
route-map set_color permit 10
 set extcommunity color 1000

mpls vpls BGP-VPLS 1000
route-map set_color
signaling bgp
ve-id 3
exit-signaling
exit-vpls

interface xe19
 mtu 9216
```

```
interface xe19.1000 switchport
 encapsulation dot1q 1000
 mtu 9216
 access-if-vpls
 mpls-vpls BGP-VPLS

router bgp 65010
bgp auto-policy-soft-reset enable
address-family l2vpn vpls
 neighbor 25.0.1.5 activate
 exit-address-family
```

R5

```
router bgp 65010
 no bgp inbound-route-filter
 bgp auto-policy-soft-reset enable
 address-family l2vpn vpls
 neighbor 25.0.1.1 activate
 neighbor 25.0.1.1 route-reflector-client
 neighbor 25.0.1.3 activate
 neighbor 25.0.1.3 route-reflector-client
 exit-address-family
 !
 exit
```

Validation

Verify the bgp vpls summary.

```
FA-RTR1#show bgp l2vpn vpls summary
BGP router identifier 25.0.1.1, local AS number 65010
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
25.0.1.5          4 65010       157       160      1      0      0  01:01:15
1

Total number of neighbors 1

Total number of Established sessions 1

FA-RTR1#show bgp l2vpn vpls detail

VPLS ID: 1000
VE-ID: 1
Discovered Peers: 1
Route-Target: 65010:1000
Local RD: 65010:1000
```

```
All Local Label Blocks:
 [LB:81280, VBO:1, VBS:64]
Mesh Peers:
 BGP Peer:25.0.1.5/32
  VC Nbr Address:25.0.1.3, RD:65010:1000, VE-ID:3
  VC Details: VC-ID:13
  Local MTU:9216, Remote MTU:9216
  Remote (LB:81280,VBO:1,VBS:64)  Local (LB:81280,VBO:1,VBS:64)
  LB sent on known VEID:Yes
  In Label:81282, Out Label:81280
  PW Status:Established
  VC Installed:Yes
  VC Signaled Time: 00:58:21
  Extended-Community Color:1000

FA-RTR1#
FA-RTR1#show mpls dep-up table
================================
Route-Node Prefix: 25.0.1.3
================================
CLIST-INFO:
    c_list-Pointer = 3579d30
    c_list-Type    = CONFIRM_NODE_FTN
    c_list-Prefix  = 25.0.1.3/32
    c_list-Count   = 1
--------------------------------
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f4627000
    Confirm-Data-Pointer  = 30b1550
    Confirm-Node-Type     = CONFIRM_VPLS_MESH_VC
    VPLS Id               =  1000
    Peer addrss           =  25.0.1.3/32
    Color                 = 1000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
FA-RTR1#
FA-RTR1#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
1000       25.0.1.3         21283        81282      ge2            81280      2/Up
1         BGP           Active     00:58:01   1000
FA-RTR1#
FA-RTR1#show mpls vpls detail
Virtual Private LAN Service Instance: BGP-VPLS, ID: 1000
 SIG-Protocol: BGP
```

```
  Route-Distinguisher :65010:1000
  Route-Target :65010:1000
  VE-ID :1
 Attachment-Circuit: UP
 Learning: Enabled
 Control-Word: Disabled
 Flow Label Status: Disabled, Direction: None, Static: No
 Group ID: 0, Configured MTU: 9216
 Description: none
 service-tpid: dot1.q
 Operating mode: Raw
 MAC Withdrawal:

 Configured interfaces:
  Interface: xe15.1000
   Status: Up
   Subinterface Match Criteria(s) :
   dot1q 1000

 Mesh Peers:
   25.0.1.3 (Type: Ethernet) (Negotiated - CW: No, FAT: No) (Up) (UpTime: 00:58:16)


FA-RTR1#
```

Verify FlexAlgo FTN entry and corresponding LSP information:

```
FA-RTR1#show mpls ftn-table 25.0.1.3/32 algorithm 128
 Primary FTN entry with FEC: 25.0.1.3/32, id: 6, row status: Active, Tunnel-Policy: N/A,
State: Installed
  CreateTime: 01:19:25, UpTime: 01:19:25, LastUpdate: N/A
  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: 8, in intf: - in label: 0 out-segment ix: 22 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 22, owner: ISIS-SR-FA, Stale: NO, refcount: 4, out intf:
ge2, out label: 21283
    Nexthop addr: 11.0.12.2        cross connect ix: 8, op code: Push

     Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 60 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 60, owner: ISIS-SR-FA, Stale: NO, refcount: 2, out intf:
ge6, out label: 21283
    Nexthop addr: 11.0.15.5        cross connect ix: 8, op code: Push

  Dependent service info (count 1):
  [VPLS_MESH_VC] VPLS mesh vpls_id 1000 peer 25.0.1.3/32 out_label 81280 in_label 81282
opcode 8 Ext-Color 1000
```

Configuration for L3VPN (VPNv4) Traffic Steering with ODN for FlexAlgo

The following configuration enables L3VPN Traffic Steering with an ODN Policy:

- This setup extends the existing configuration from the Flex Algorithm for ISIS chapter to support L3VPN services over Flex-Algo 128.
- ODN policy automates the steering of L3VPN traffic, ensuring it follows optimized paths based on the constraints and metrics defined for Flex-Algo 128.
- This approach enables seamless integration of L3VPN over an SR-based network while providing traffic engineering flexibility and efficient resource utilization.

1. Configure BGP On RTR1 and RTR3 which are PE routers and RTR5(RR).

    1. Set up BGP in Address-Family VPNv4 on RTR1 and RTR2 as below.

```
RTR1(config)#router bgp 65010
RTR1(config-router)# bgp auto-policy-soft-reset enable
RTR1(config-router)# neighbor 25.0.1.5 remote-as 65010
RTR1(config-router)# neighbor 25.0.1.5 update-source loopback2
RTR1(config-router)# end
RTR1(config-router)# address-family vpnv4 unicast
RTR1(config-router-af)# neighbor 25.0.1.5 activate
RTR1(config-router-af)# exit-address-family
RTR1(config-router)# end
RTR1(config-router)# exit

RTR3(config)#router bgp 65010
RTR3(config-router)# bgp auto-policy-soft-reset enable
RTR3(config-router)# neighbor 25.0.1.5 remote-as 65010
RTR3(config-router)# neighbor 25.0.1.5 update-source loopback2
RTR3(config-router)# end
RTR3(config-router)# address-family vpnv4 unicast
RTR3(config-router-af)# neighbor 25.0.1.5 activate
RTR3(config-router-af)# exit-address-family
RTR3(config-router)# end
RTR3(config-router)# exit
```

    2. Set up BGP in Address-Family VPNv4 on RTR3 which acts as RR below:

```
RTR5(config)#router bgp 65010
RTR5(config-router)# bgp auto-policy-soft-reset enable
RTR5(config-router)# no bgp inbound-route-filter
RTR5(config-router)# neighbor 25.0.1.1 remote-as 65010
RTR5(config-router)# neighbor 25.0.1.1 update-source loopback2
RTR5(config-router)# neighbor 25.0.1.3 remote-as 65010
RTR5(config-router)# neighbor 25.0.1.3 update-source loopback2
RTR5(config-router)# end
RTR5(config-router)# address-family vpnv4 unicast
RTR5(config-router-af)# neighbor 25.0.1.1 activate
RTR5(config-router-af)# neighbor 25.0.1.1 route-reflector-client
RTR5(config-router-af)# neighbor 25.0.1.3 activate
RTR5(config-router-af)# neighbor 25.0.1.3 route-reflector-client
RTR5(config-router-af)# exit-address-family
RTR5(config-router)# end
RTR5(config-router)# exit
```

2. Configure VRF on PE Routers (RTR1 and RTR3).

```
RTR1(config)#ip vrf vrf2000
RTR1(config-vrf)# rd 10:2000
RTR1(config-vrf)# route-target both 10:2000


RTR3(config)#ip vrf vrf2000
RTR3(config-vrf)# rd 10:2000
RTR3(config-vrf)# route-target both 10:2000
```

3. Enable BGP for VRF and Configure ACCESS-IF:

```
RTR1(config)#int xe15
RTR1(config-if)# mtu 9216
RTR1(config-if)#exit

RTR1(config)#
RTR1(config-if)#interface xe15.2000
RTR1(config-if)# encapsulation dot1q 2000
RTR1(config-if)# ip vrf forwarding vrf2000
RTR1(config-if)# ip address 100.1.1.1/24
RTR1(config-if)# mtu 9216
RTR1(config-if)#exit

RTR1(config)#
RTR1(config)#router bgp 65010
RTR1(config-router)# address-family ipv4 vrf vrf2000
RTR1(config-router-af)# redistribute connected
RTR1(config-router-af)# neighbor 100.1.1.2 remote-as 100
RTR1(config-router-af)# neighbor 100.1.1.2 activate
RTR1(config-router-af)# exit-address-family
RTR1(config-router)#end
RTR1(config-router)#commit


RTR3(config)#
RTR3(config)#int xe19
RTR3(config-if)# mtu 9216
RTR3(config-if)#exit

RTR3(config)#
RTR3(config-if)#interface xe19.2000
RTR3(config-if)# encapsulation dot1q 2000
RTR3(config-if)# ip vrf forwarding vrf2000
RTR3(config-if)# ip address 200.1.1.1/24
RTR3(config-if)# mtu 9216
RTR3(config-if)#exit

RTR3(config)#
RTR3(config)#router bgp 65010
RTR3(config-router)# address-family ipv4 vrf vrf2000
RTR3(config-router-af)# redistribute connected
RTR3(config-router-af)# neighbor 200.1.1.2 remote-as 200
```

```
RTR3(config-router-af)# neighbor 200.1.1.2 activate
RTR3(config-router-af)# exit-address-family
RTR3(config-router)#end
RTR3(config-router)#commit
```

4. Configure the egress side to advertise the color per VRF under address-family using a route-map.

```
RTR3(config)#
RTR3(config)#route-map set_color_vrf2000 permit 10
RTR3(config-route-map)#set extcommunity color 2000
RTR3(config-if)#exit

RTR3(config)#router bgp 65010
RTR3(config-router)# address-family ipv4 vrf vrf2000
RTR3(config-router-af)# redistribute connected route-map set_color_vrf2000
RTR3(config-router-af)# neighbor 200.1.1.2 route-map set_color_vrf2000 in
RTR3(config-router-af)# exit-address-family
RTR3(config-router)#
```

5. Enable ODN Policy for Traffic Steering on RTR1 for L3VPN routes coming from RTR3:

```
RTR1(config)#
RTR1(config)#segment-routing
RTR1(config-sr)# traffic-engineering
RTR1(config-sr-te)#  on-demand-nexthop 2000
RTR1(config-sr-odn)#   flex-algo 128
RTR1(config-sr-odn)#  exit-sr-odn
RTR1(config-sr-te)#  !
RTR1(config-sr-te)# exit-te
RTR1(config-sr)#commit
```

Configuration Snapshot

RTR1

```
router bgp 65010
bgp auto-policy-soft-reset enable
neighbor 25.0.1.5 remote-as 65010
neighbor 25.0.1.5 update-source loopback2
!
address-family vpnv4 unicast
 neighbor 25.0.1.5 activate
 exit-address-family
!
exit

ip vrf vrf2000
 rd 10:2000
 route-target both 10:2000

int xe15
 mtu 9216

interface xe15.2000
 encapsulation dot1q 2000
```

```
 ip vrf forwarding vrf2000
 ip address 100.1.1.1/24
 mtu 9216

!
router bgp 65010
 address-family ipv4 vrf vrf2000
 redistribute connected
 neighbor 100.1.1.2 remote-as 100
 neighbor 100.1.1.2 activate
 exit-address-family
!

segment-routing
 traffic-engineering
  on-demand-nexthop 2000
   flex-algo 128
  exit-sr-odn
  !
 exit-te
```

RTR3

```
router bgp 65010
bgp auto-policy-soft-reset enable
neighbor 25.0.1.5 remote-as 65010
neighbor 25.0.1.5 update-source loopback2
!
address-family vpnv4 unicast
 neighbor 25.0.1.5 activate
 exit-address-family
!
exit

ip vrf vrf2000
 rd 10:2000
 route-target both 10:2000

int xe19
 mtu 9216

interface xe19.2000
 encapsulation dot1q 2000
 ip vrf forwarding vrf2000
 ip address 200.1.1.1/24
 mtu 9216

!
router bgp 65010
 address-family ipv4 vrf vrf2000
 redistribute connected
 neighbor 200.1.1.2 remote-as 200
```

```
 neighbor 200.1.1.2 activate
 exit-address-family
!

route-map set_color_vrf2000 permit 10
set extcommunity color 2000

router bgp 65010
 address-family ipv4 vrf vrf2000
 redistribute connected route-map set_color_vrf2000
 neighbor 200.1.1.2 remote-as 200
 neighbor 200.1.1.2 activate
 neighbor 200.1.1.2 route-map set_color_vrf2000 in
 exit-address-family
```

RTR5

```
router bgp 65010
 bgp auto-policy-soft-reset enable
 no bgp inbound-route-filter
 neighbor 25.0.1.1 remote-as 65010
 neighbor 25.0.1.1 update-source loopback2
 neighbor 25.0.1.3 remote-as 65010
 neighbor 25.0.1.3 update-source loopback2
 !
 address-family vpnv4 unicast
 neighbor 25.0.1.1 activate
 neighbor 25.0.1.1 route-reflector-client
 neighbor 25.0.1.3 activate
 neighbor 25.0.1.3 route-reflector-client
 exit-address-family
 !
exit
```

Validation

Ensure the VPNv4 routes are properly advertised and received across the BGP sessions. Check for correct route attributes and steering behaviors.:

```
FA-RTR1#show ip bgp vpnv4 all summary
BGP router identifier 25.0.1.1, local AS number 65010
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   Desc
25.0.1.5          4 65010        526        538       5      0      0  01:40:52
11

Total number of neighbors 1

Total number of Established sessions 1
```

```
BGP VRF vrf2000 Route Distinguisher: 10:2000
BGP table version is 2
3 BGP AS-PATH entries
0 BGP community entries

Neighbor          V    AS     MsgRcv     MsgSen  TblVer    InQ   OutQ   Up/Down  State/
PfxRcd   Desc
100.1.1.2         4   100         41         49       2      0      0  00:15:47
10

Total number of neighbors 1

Total number of Established sessions 1
FA-RTR1#sh ip bgp vpnv4 all
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

     Network          Next Hop            Metric    LocPrf     Weight Path
Route Distinguisher: 10:2000 (Default for VRF vrf2000)
*> l 100.1.1.0/24     0.0.0.0                  0        100       32768  ?        -
*> l 101.0.0.0/24     100.1.1.2                0        100       0   100 i        -
*> l 101.0.1.0/24     100.1.1.2                0        100       0   100 i        -
*> l 101.0.2.0/24     100.1.1.2                0        100       0   100 i        -
*> l 101.0.3.0/24     100.1.1.2                0        100       0   100 i        -
*> l 101.0.4.0/24     100.1.1.2                0        100       0   100 i        -
*> l 101.0.5.0/24     100.1.1.2                0        100       0   100 i        -
*> l 101.0.6.0/24     100.1.1.2                0        100       0   100 i        -
*> l 101.0.7.0/24     100.1.1.2                0        100       0   100 i        -
*> l 101.0.8.0/24     100.1.1.2                0        100       0   100 i        -
*> l 101.0.9.0/24     100.1.1.2                0        100       0   100 i        -
*>i  200.1.1.0        25.0.1.3                 0        100       0    ?       2000
*>i  201.0.0.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.1.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.2.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.3.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.4.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.5.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.6.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.7.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.8.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.9.0        25.0.1.3                 0        100       0   200 i       2000
 Announced routes count = 11
 Accepted routes count = 11
Route Distinguisher: 10:2000
*>i  200.1.1.0        25.0.1.3                 0        100       0    ?       2000
*>i  201.0.0.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.1.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.2.0        25.0.1.3                 0        100       0   200 i       2000
*>i  201.0.3.0        25.0.1.3                 0        100       0   200 i       2000
```

```
*>i  201.0.4.0        25.0.1.3             0        100       0   200 i       2000
*>i  201.0.5.0        25.0.1.3             0        100       0   200 i       2000
*>i  201.0.6.0        25.0.1.3             0        100       0   200 i       2000
*>i  201.0.7.0        25.0.1.3             0        100       0   200 i       2000
*>i  201.0.8.0        25.0.1.3             0        100       0   200 i       2000
*>i  201.0.9.0        25.0.1.3             0        100       0   200 i       2000
 Announced routes count = 0
 Accepted routes count = 11
FA-RTR1#
FA-RTR1#sh ip route vrf vrf2000
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 "vrf2000"
C            100.1.1.0/24 is directly connected, xe15.2000, installed 00:35:21, last
update 00:35:21 ago
B            101.0.0.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
B            101.0.1.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
B            101.0.2.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
B            101.0.3.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
B            101.0.4.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
B            101.0.5.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
B            101.0.6.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
B            101.0.7.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
B            101.0.8.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
B            101.0.9.0/24 [20/0] via 100.1.1.2, xe15.2000, installed 00:16:14, last
update 00:16:14 ago
C            127.0.0.0/8 is directly connected, lo.vrf2000, installed 01:38:23, last
update 01:38:23 ago

Gateway of last resort is not set
```

Ensure that MPLS label switching follows the path determined by FlexAlgo and ODN.

```
FA-RTR1#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 LDP ECMP or SR ECMP
Ext-Color - Extended-community color advertised by BGP
```

```
Code    FEC               FTN-ID VRF-ID    Nhlfe-ID    Pri   Out-Label    Out-Intf
Nexthop         UpTime    Ext-Color
   B>   200.1.1.0/24       1        2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.0.0/24       2        2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.1.0/24       3        2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.2.0/24       4        2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.3.0/24       5        2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.4.0/24       6        2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.5.0/24       7        2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.6.0/24       8        2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.7.0/24       9        2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.8.0/24       10       2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
   B>   201.0.9.0/24       11       2      4           Yes   81344        -
25.0.1.3        00:06:36    2000
FA-RTR1#
FA-RTR1#show mpls dep-up table
================================
Route-Node Prefix: 25.0.1.3
================================
CLIST-INFO:
    c_list-Pointer = 30b33e0
    c_list-Type    = CONFIRM_NODE_FTN
    c_list-Prefix  = 25.0.1.3/32
    c_list-Count   = 12
--------------------------------
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f4627060
    Confirm-Data-Pointer  = ec5d4400
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.0.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f4627090
    Confirm-Data-Pointer  = ec5d4740
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.1.0/24
    Color                 = 2000
```

```
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f46270c0
    Confirm-Data-Pointer  = ec5d4a80
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.2.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f46270f0
    Confirm-Data-Pointer  = ec5d4dc0
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.3.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f4627120
    Confirm-Data-Pointer  = ec5d5100
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.4.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f4627150
    Confirm-Data-Pointer  = ec5d5440
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.5.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
```

```
    Confirm-Node-Pointer  = f4627180
    Confirm-Data-Pointer  = ec5d5780
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.6.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f46271b0
    Confirm-Data-Pointer  = ec5d5ac0
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.7.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f46271e0
    Confirm-Data-Pointer  = ec5d5e00
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.8.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f4627210
    Confirm-Data-Pointer  = ec5d6140
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 201.0.9.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = f4627030
    Confirm-Data-Pointer  = ec5d40c0
    Confirm-Node-Type     = CONFIRM_VRF
    Fec-Prefix            = 200.1.1.0/24
    Color                 = 2000
      Parent-FTN-Pointer  = ec5d2040
```

```
      Parent-FTN-Index    = 6
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 128

FA-RTR1#
FA-RTR1#
FA-RTR1#show mpls forwarding-table 25.0.1.3/32 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>   25.0.1.3/32       6       61       -         (e)   -          -         -
-              128        03:52:48
                                    22       0          Yes   21283        ge2
No        11.0.12.2       -          -
                                    60       0          Yes   21283        ge6
No        11.0.15.5       -          -
FA-RTR1#show mpls ftn-table 25.0.1.3/32 algorithm 128
 Primary FTN entry with FEC: 25.0.1.3/32, id: 6, row status: Active, Tunnel-Policy: N/A,
State: Installed
  CreateTime: 03:52:56, UpTime: 03:52:56, LastUpdate: N/A
  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: 8, in intf: - in label: 0 out-segment ix: 22 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 22, owner: ISIS-SR-FA, Stale: NO, refcount: 4, out intf:
ge2, out label: 21283
    Nexthop addr: 11.0.12.2        cross connect ix: 8, op code: Push

     Cross connect ix: 8, in intf: - in label: 0 out-segment ix: 60 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 60, owner: ISIS-SR-FA, Stale: NO, refcount: 2, out intf:
ge6, out label: 21283
    Nexthop addr: 11.0.15.5        cross connect ix: 8, op code: Push

  Dependent service info (count 12):
  [CONFIRM_VRF] ftn_ix 6 owner BGP prefix 201.0.4.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
  [CONFIRM_VRF] ftn_ix 3 owner BGP prefix 201.0.1.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
  [CONFIRM_VRF] ftn_ix 1 owner BGP prefix 200.1.1.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
  [CONFIRM_VRF] ftn_ix 2 owner BGP prefix 201.0.0.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
  [CONFIRM_VRF] ftn_ix 4 owner BGP prefix 201.0.2.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
  [CONFIRM_VRF] ftn_ix 5 owner BGP prefix 201.0.3.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
  [CONFIRM_VRF] ftn_ix 9 owner BGP prefix 201.0.7.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
```

```
[CONFIRM_VRF] ftn_ix 7 owner BGP prefix 201.0.5.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
[CONFIRM_VRF] ftn_ix 8 owner BGP prefix 201.0.6.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
[CONFIRM_VRF] ftn_ix 10 owner BGP prefix 201.0.8.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
[CONFIRM_VRF] ftn_ix 11 owner BGP prefix 201.0.9.0/24 nhlfe_ix 4 vrf 2 Ext-Color 2000
```

Configuration for EVPN Traffic Steering with ODN for FlexAlgo

The following configuration enables EVPN-ELINE Traffic Steering with an ODN Policy:

- This setup extends the existing configuration from the Flex Algorithm for ISIS chapter to support EVPN ELINE services over Flex-Algo 129.
- The ODN policy dynamically steers EVPN ELINE traffic, ensuring that it follows the most optimal paths based on the constraints and metrics defined for Flex-Algo 129.
- This approach facilitates the seamless integration of EVPN ELINE over an SR-based network, enhancing traffic engineering flexibility while optimizing resource utilization.

1. Configure BGP On RTR1 and RTR3 which are PE routers and RTR5(RR).

    1. Set up BGP in Address-Family L2VPN EVPN on RTR1 and RTR2 as below.

```
RTR1(config)#router bgp 65010
RTR1(config-router)# bgp auto-policy-soft-reset enable
RTR1(config-router)# neighbor 25.0.1.5 remote-as 65010
RTR1(config-router)# neighbor 25.0.1.5 update-source loopback2
RTR1(config-router)# end
RTR1(config-router)# address-family l2vpn evpn
RTR1(config-router-af)# neighbor 25.0.1.5 activate
RTR1(config-router-af)# exit-address-family
RTR1(config-router)# end
RTR1(config-router)# exit
```

```
RTR3(config)#router bgp 65010
RTR3(config-router)# bgp auto-policy-soft-reset enable
RTR3(config-router)# neighbor 25.0.1.5 remote-as 65010
RTR3(config-router)# neighbor 25.0.1.5 update-source loopback2
RTR3(config-router)# end
RTR3(config-router)# address-family l2vpn evpn
RTR3(config-router-af)# neighbor 25.0.1.5 activate
RTR3(config-router-af)# exit-address-family
RTR3(config-router)# end
RTR3(config-router)# exit
```

    2. Set up BGP in Address-Family L2VPN EVPN on RTR3 which acts as RR below:

```
RTR5(config)#router bgp 65010
RTR5(config-router)# bgp auto-policy-soft-reset enable
RTR5(config-router)# no bgp inbound-route-filter
RTR5(config-router)# neighbor 25.0.1.1 remote-as 65010
RTR5(config-router)# neighbor 25.0.1.1 update-source loopback2
RTR5(config-router)# neighbor 25.0.1.3 remote-as 65010
RTR5(config-router)# neighbor 25.0.1.3 update-source loopback2
RTR5(config-router)# !
RTR5(config-router)# address-family l2vpn evpn
RTR5(config-router-af)# neighbor 25.0.1.1 activate
RTR5(config-router-af)# neighbor 25.0.1.1 route-reflector-client
```

```
RTR5(config-router-af)# neighbor 25.0.1.3 activate
RTR5(config-router-af)# neighbor 25.0.1.3 route-reflector-client
RTR5(config-router-af)# exit-address-family
RTR5(config-router)# !
RTR5(config-router)# exit
```

2. Configure MAC-VRF on PE Routers (RTR1 and RTR3):

```
RTR1(config)#ip vrf vrf2000
RTR1(config-vrf)# rd 10:2000
RTR1(config-vrf)# route-target both 10:2000

RTR3(config)#ip vrf vrf2000
RTR3(config-vrf)# rd 10:2000
RTR3(config-vrf)# route-target both 10:2000
```

3. Configure EVPN ELINE instance:

```
RTR1(config)#evpn mpls enable
RTR1(config)#evpn mpls vtep-ip-global 25.0.1.1
RTR1(config)#!
RTR1(config)#evpn mpls id 3000 xconnect target-mpls-id 3001
RTR1(config-evpn-mpls)# host-reachability-protocol evpn-bgp ELINE_SH
RTR1(config-evpn-mpls)#!


RTR3(config)#evpn mpls enable
RTR3(config)#evpn mpls vtep-ip-global 25.0.1.3
RTR3(config)#!
RTR3(config)#evpn mpls id 3001 xconnect target-mpls-id 3000
RTR3(config-evpn-mpls)# host-reachability-protocol evpn-bgp ELINE_SH
RTR3(config-evpn-mpls)#!
```

4. Configure the egress side to advertise the color per VRF under address-family using a route-map.

```
RTR3(config)#
RTR3(config)#route-map set_color_ELINE3000 permit 10
RTR3(config-route-map)#set extcommunity color 3000
RTR3(config-if)#exit

RTR3(config)#evpn mpls id 3001 xconnect target-mpls-id 3000
RTR3(config-evpn-mpls)# route-map set_color_ELINE3000
RTR3(config-evpn-mpls)#end
```

5. Enable ODN Policy for Traffic Steering on RTR1 for EVPN routes coming from RTR3

```
RTR1(config)#
RTR1(config)#segment-routing
RTR1(config-sr)# traffic-engineering
RTR1(config-sr-te)# on-demand-nexthop 3000
RTR1(config-sr-odn)# flex-algo 129
RTR1(config-sr-odn)# exit-sr-odn
RTR1(config-sr-te)# end
RTR1(config-sr-te)# exit-te
RTR1(config-sr)#commit
```

Configuration Snapshot

RTR1

```
router bgp 65010
bgp auto-policy-soft-reset enable
neighbor 25.0.1.5 remote-as 65010
neighbor 25.0.1.5 update-source loopback2
!
address-family l2vpn evpn
 neighbor 25.0.1.5 activate
 exit-address-family
!
exit

mac vrf ELINE_SH
 rd 25.0.1.1:3000
 route-target both evpn-auto-rt

evpn mpls enable

evpn mpls vtep-ip-global 25.0.1.1
!
evpn mpls id 3000 xconnect target-mpls-id 3001
 host-reachability-protocol evpn-bgp ELINE_SH
!
interface xe15
mtu 9216
!
interface xe15.3000 switchport
 encapsulation dot1q 3000
 access-if-evpn
  map vpn-id 3000
```

RTR3

```
router bgp 65010
bgp auto-policy-soft-reset enable
neighbor 25.0.1.5 remote-as 65010
neighbor 25.0.1.5 update-source loopback2
!
address-family l2vpn evpn
 neighbor 25.0.1.5 activate
 exit-address-family
!
!
exit

mac vrf ELINE_SH
 rd 25.0.1.3:3000
 route-target both evpn-auto-rt

evpn mpls enable

evpn mpls vtep-ip-global 25.0.1.3
```

```
!
evpn mpls id 3001 xconnect target-mpls-id 3000
 host-reachability-protocol evpn-bgp ELINE_SH
!
interface xe19
mtu 9216
!
interface xe19.3000 switchport
 encapsulation dot1q 3000
 mtu 9216
 access-if-evpn
  map vpn-id 3001

route-map set_color_ELINE3000 permit 10
set extcommunity color 3000
exit

evpn mpls id 3001 xconnect target-mpls-id 3000
 route-map set_color_ELINE3000
```

RTR5

```
router bgp 65010
 bgp auto-policy-soft-reset enable
 no bgp inbound-route-filter
 neighbor 25.0.1.1 remote-as 65010
 neighbor 25.0.1.1 update-source loopback2
 neighbor 25.0.1.3 remote-as 65010
 neighbor 25.0.1.3 update-source loopback2
 !
 address-family l2vpn evpn
 neighbor 25.0.1.1 activate
 neighbor 25.0.1.1 route-reflector-client
 neighbor 25.0.1.3 activate
 neighbor 25.0.1.3 route-reflector-client
 exit-address-family
 !
exit
```

Validation

Verify BGP EVPN neighbor state and received prefixes:

```
FA-RTR1#show bgp l2vpn evpn summary
BGP router identifier 25.0.1.1, local AS number 65010
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     AD  MACIP  MCAST    ESI  PREFIX-ROUTE  Desc
```

```
25.0.1.5          4 65010       123       113      8      0      0  00:44:41
1      1      0      0      0      0

Total number of neighbors 1

Total number of Established sessions 1
FA-RTR1#sh bgp l2vpn evpn
BGP table version is 8, local router ID is 25.0.1.1
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

[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[25.0.1.1:3000] VRF[ELINE_SH]:
*>   [1]:[0]:[3000]:[81920]
                       25.0.1.1             0        100       32768  i        -  -----
-----      MPLS
* i  [1]:[0]:[3001]:[81280]
                    25.0.1.3           0       100      0    i      3000  25.0.1.5
MPLS

RD[25.0.1.3:3000]
*>i  [1]:[0]:[3001]:[81280]
                    25.0.1.3           0       100      0    i      3000  25.0.1.5
MPLS
```

Verify dependency resolution for EVPN service steering over FlexAlgo:

```
FA-RTR1#show mpls dep-up table
================================
Route-Node Prefix: 25.0.1.3
================================
CLIST-INFO:
    c_list-Pointer = 27d34d0
    c_list-Type    = CONFIRM_NODE_FTN
    c_list-Prefix  = 25.0.1.3/32
    c_list-Count   = 1
--------------------------------
CONFIRM-NODE INFO:
    Confirm-Node-Pointer  = b6402210
    Confirm-Data-Pointer  = 27d3e30
    Confirm-Node-Type     = CONFIRM_EVPN
```

```
    EVPN-ID               = 3000
    Destination Peer      = 25.0.1.3
    Color                 = 3000
      Parent-FTN-Pointer  = ae7c0700
      Parent-FTN-Index    = 13
      Parent-FTN-Name     = N/A
      Parent-FTN-Owner    = ISIS-SR-FA
      Parent-FTN-Algo-Num = 129
```

Verify EVPN MPLS tunnel label allocation and underlay path:

```
FA-RTR1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source           Destination      Status        Up/Down       Update        local-evpn-
id remote-evpn-id
================================================================================
=================
25.0.1.1         25.0.1.3         Installed     00:05:12      00:05:12      3000
3001

Total number of entries are 1
```

Verify EVPN MPLS tunnel label allocation and underlay path:

```
FA-RTR1#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
================+===========+=========+=========+===============+=========+=========+==
==========+=========+============+=========
 25.0.1.3       Installed   3000     3001     --            81920    81280    --
57        ge2           21293

Total number of entries are 1
```

Verify EVPN cross-connect status and MTU configuration:

```
FA-RTR1#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
================================ ============
================================================================================
3000         ----          9216  3001         xe15.3000    --- Single Homed Port ---
25.0.1.3        9216  AC-NW  NW-SET

Total number of entries are 1
FA-RTR1#
```

Verify FlexAlgo FTN entry and corresponding LSP information:

```
FA-RTR1#show mpls ftn-table 25.0.1.3/32 algorithm 129
 Primary FTN entry with FEC: 25.0.1.3/32, id: 13, row status: Active, Tunnel-Policy: N/
A, State: Installed
  CreateTime: 00:45:21, UpTime: 00:45:21, LastUpdate: N/A
  Owner: ISIS-SR-FA, distance: 115, Action-type: Redirect to LSP, Exp-bits: 0x0,
Incoming DSCP: none, Algorithm Number:129
  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: 56 refcount: 1
      Owner: ISIS-SR-FA, Persistent: No, Admin Status: Up, Oper Status: Up
      State: Active
       Out-segment with ix: 56, owner: ISIS-SR-FA, Stale: NO, refcount: 2, out intf:
ge2, out label: 21293
    Nexthop addr: 11.0.12.2        cross connect ix: 7, op code: Push
```

CHAPTER 5 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.

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.

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 connec-tivity. 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 net-work. 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.

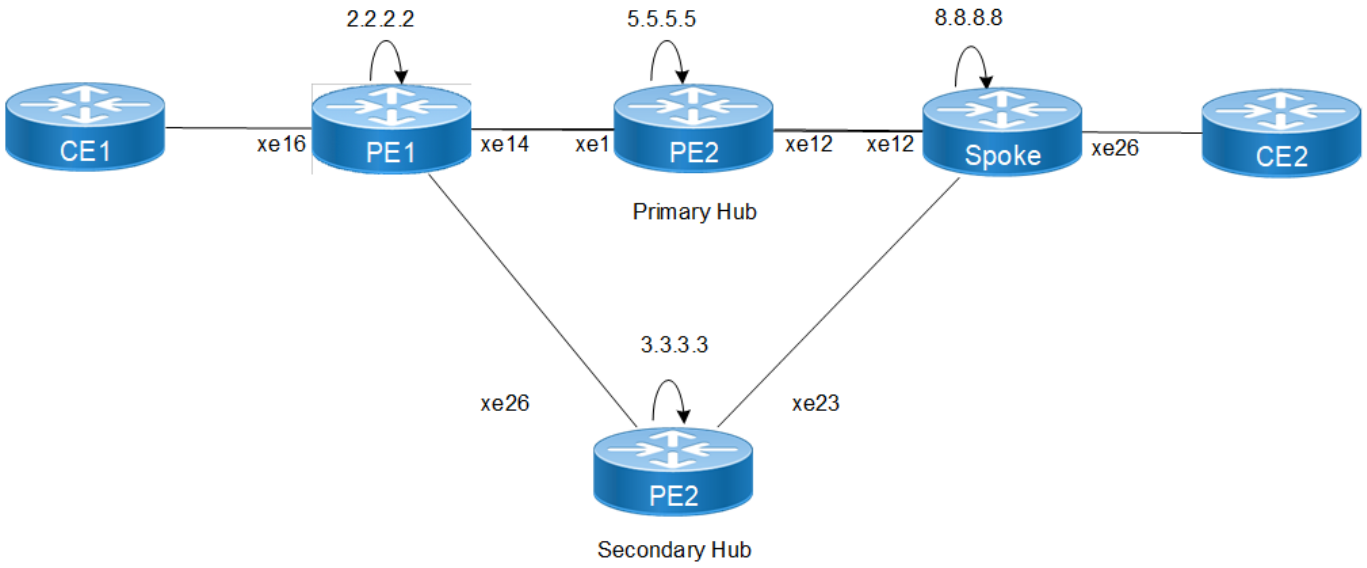

MAC Withdrawal for VPLS/H-VPLS Topology

Configuring H-VPLS

Configure PE1 router as follows:

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)#
```

Configuring the H-VPLS MAC-Withdrawal 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)#
```

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)#
```

Configuring 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. 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)#
```

Configuring 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)#
```

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)#
```

Running Configuration on PE1 Router:

```
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
```

Running Configuration on PE2 Router:

```
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
```

Running Configuration on PE3 Router:

```
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
```

Running Configuration on 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
mac-withdrawal flush-on-spoke-vc-standby-activation
 exit-vpls
!
interface xe26.2000 switchport
 access-if-vpls
    mpls-vpls vpls2000
```

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)
```

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
   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)
```

CLI Commands

The MAC Withdrawal for VPLS/H-VPLS 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

Explain or describe the example.

```
#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
```

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
```

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

CHAPTER 6 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 connec-tivity. 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 net-work. 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.

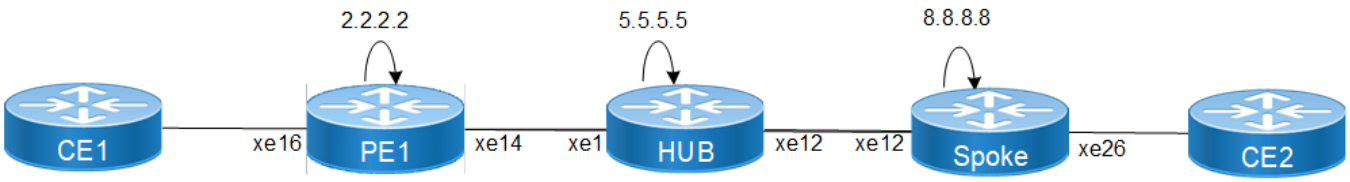

MAC Move Protection Topology Diagram

Configuring MAC Move Protection

Configure PE1 router as follows:

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)#
```

Configure the 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)#
```

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)#
```

Configure Spoke Router as follows:

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 on PE1 Router:

```
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
```

Running Configuration on Hub Router:

```
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
```

Running Configuration on Spoke Router:

```
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: dot1.q
 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: dot1.q
 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: dot1.q
```

```
 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

`name`	Specifies the name of VPLS instance.

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

CHAPTER 7 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 Protection - 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 connec-tivity. 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 net-work. 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.

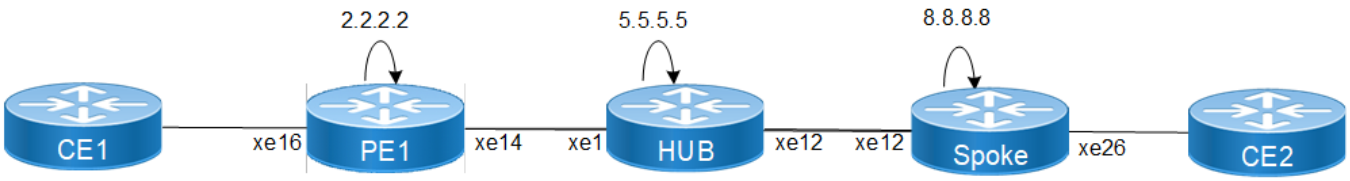

MAC Limit for H-VPLS Topology Diagram

Configuring MAC-Limit

Configure PE1 router as follows:

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
```

Configure the MAC-LIMIT 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)#
```

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)#
```

Configure the MAC-LIMIT 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)#
```

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 on PE1 Router:

```
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
```

Running Configuration on Hub Router:

```
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 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
 exit-vpls
```

Running Configuration on Spoke Router:

```
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       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                      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                      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

CHAPTER 8 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.

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 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.
- VRRP over the BVI interface is not supported.
- TWAMP support along with QoS shaping and Queuing for BVI is planned for future releases.
- 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).

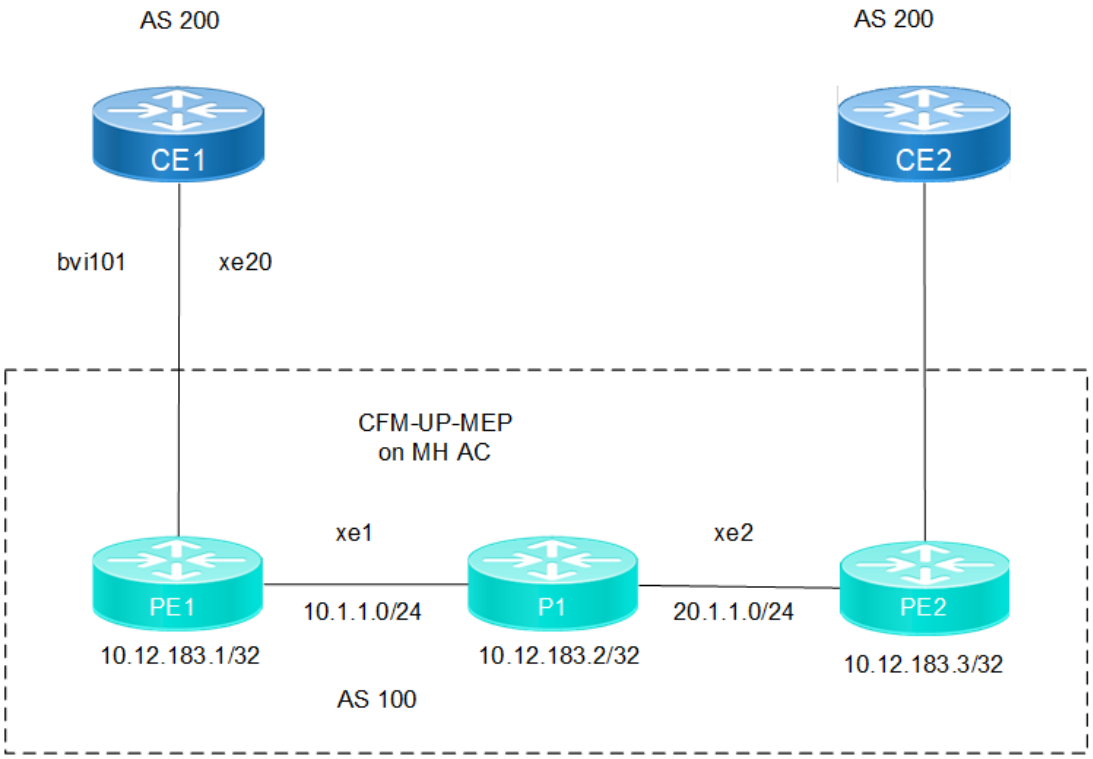

BVI Interface Topology

CE1

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. Configure BGP 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
```

PE1

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 (Bridge Virtual Interface) 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-f)# 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 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
```

PE2

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
```

CE2

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 101.0.1.1 remote-as 100
CE2(config-router)#address-family ipv4 unicast
CE2(config-router-af)#neighbor 101.0.1.1 activate
CE2(config-router-af)#exit-address-family
CE2(config-router)#exit
```

Running configuration on CE1 router is as follows:

```
#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
```

Running configuration on PE1 router is as follows:

```
#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 bvi101
 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 xe1
 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
```

Running configuration on P1 router is as follows:

```
#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
```

Running configuration on PE2 router is as follows:

```
#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
```

Running configuration on CE2 router is as follows:

```
#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

CHAPTER 9 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.

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.

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 net-work. 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 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
!
```

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.

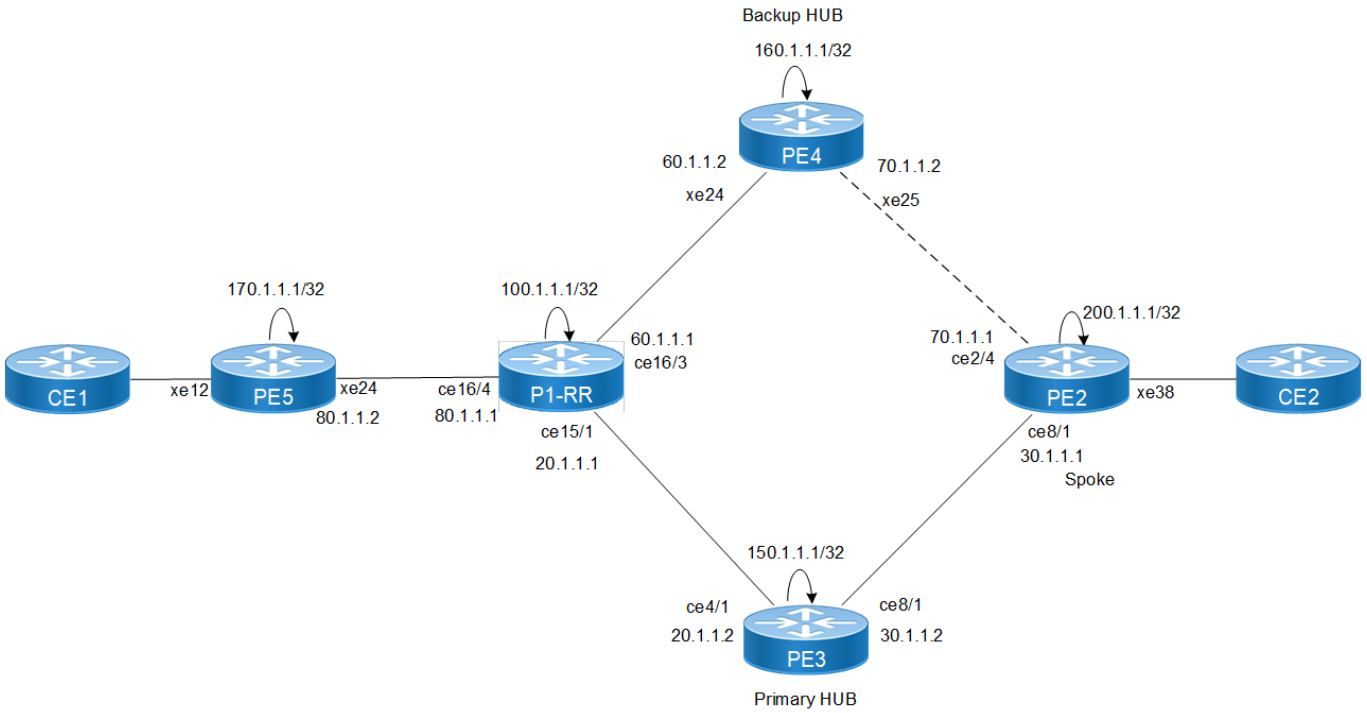

BGP-A-D LDP VPLS Topology Diagram

Configuring BGP-A-D LDP VPLS

Configure PE5 router as follows:

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)#signaling ldp
PE5(config-vpls-sig)#bgp-auto-discovery
PE5(config-vpls-ldp-sig-bgp-ad)#l2vpn-id 200:1001
PE5(config-vpls-ldp-sig-bgp-ad)#rd 10.10.10.10:1001
PE5(config-vpls-ldp-sig-bgp-ad)#route-target both 2:100
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
```

Configure P1 router as follows:

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
```

Configure PE3 router as follows:

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-vc vc1
PE3(config-vpls-spoke)#exit-spoke
PE3(config-vpls)#signaling ldp
PE3(config-vpls-sig)#bgp-auto-discovery
PE3(config-vpls-ldp-sig-bgp-ad)#l2vpn-id 200:1001
PE3(config-vpls-ldp-sig-bgp-ad)#rd 10.10.10.10:1001
PE3(config-vpls-ldp-sig-bgp-ad)#route-target both 2:100
PE3(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery
PE3(config-vpls-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
```

Configure PE4 router as follows:

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-vc vc1001
PE4(config-vpls-spoke)#exit-spoke
PE4(config-vpls)#signaling ldp
PE4(config-vpls-sig)#bgp-auto-discovery
PE4(config-vpls-ldp-sig-bgp-ad)#l2vpn-id 200:1001
PE4(config-vpls-ldp-sig-bgp-ad)#rd 10.10.10.10:1001
PE4(config-vpls-ldp-sig-bgp-ad)#route-target both 2:100
PE4(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery
PE4(config-vpls-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
```

Configure PE2 router as follows:

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 on PE1 Router:

LDP:

```
router ldp
 router-id 2.2.2.2
 pw-status-tlv
 transport-address ipv4 2.2.2.2
!
!
interface xe14
     enable-ldp ipv4
!
```

VPLS:

```
mpls vpls vpls2000 2000
signaling ldp
  bgp-auto-discovery
    rd 100.200.100.200:200
    route-target both 64000:1
    l2vpn-id 64000:2000
    exit-bgp-auto-discovery
        exit-signaling
```

BGP:

```
router bgp 64000
 bgp router-id 2.2.2.2
 neighbor 3.3.3.3 remote-as 64000
 neighbor 5.5.5.5 remote-as 64000
 neighbor 3.3.3.3 update-source lo
 neighbor 5.5.5.5 update-source lo
 !
 address-family l2vpn vpls
 neighbor 3.3.3.3 activate
 neighbor 5.5.5.5 activate
```

Running Configuration on PE2 Router:

```
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
!
```

Running Configuration on PE3 Router:

```
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 vpls vpls2 2
 vpls-vc vc1
 exit-spoke
 signaling ldp
  bgp-auto-discovery
    l2vpn-id 200:1001
    rd 10.10.10.10:1001
    route-target both 2:100
    exit-bgp-auto-discovery
  exit-signaling
 exit-vpls
!
interface ce4/2.2 switchport
 encapsulation dot1q 2
  access-if-vpls
    mpls-vpls vpls2
  exit
```

```
 exit
!
end
```

Running Configuration on P1 Router:

```
router ldp
 router-id 100.1.1.1
 transport-address ipv4 100.1.1.1
!
interface ce15/1
 ip address 20.1.1.1/24
 label-switching
 mpls ldp-igp sync ospf
 ip ospf network point-to-point
 enable-ldp ipv4
!
interface ce16/3
 ip address 60.1.1.1/24
 label-switching
 mpls ldp-igp sync ospf
 ip ospf network point-to-point
 enable-ldp ipv4
!
interface ce16/4
 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
!
```

Running Configuration on PE4 Router:

```
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-vc vc1001
 exit-spoke
 signaling ldp
  bgp-auto-discovery
    l2vpn-id 200:1001
    rd 10.10.10.10:1001
    route-target both 2:100
    exit-bgp-auto-discovery
  exit-signaling
 exit-vpls
!
interface xe12.2 switchport
 encapsulation dot1q 2
  access-if-vpls
    mpls-vpls vpls2
  exit
 exit
!
end
```

Running Configuration on PE5 Router:

```
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
 signaling ldp
  bgp-auto-discovery
    l2vpn-id 200:1001
    rd 10.10.10.10:1001
    route-target both 2:100
    exit-bgp-auto-discovery
  exit-signaling
 exit-vpls
!
interface xe12.2 switchport
 encapsulation dot1q 2
  access-if-vpls
    mpls-vpls vpls2
  exit
 exit
!
```

Validation

Validate the show output after configuration as shown below.

```
PE3-Pri-HUB#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
2          160.1.1.1         24322         26241      ce4/1          26241      2/Up
0         LDP           Active     00:02:28      -
2          170.1.1.1         25612         26242      ce4/1          26240      2/Up
0         LDP           Active     00:00:31      -
9600-PE3-Pri-HUB#
9600-PE3-Pri-HUB#
9600-PE3-Pri-HUB#sh mpls vpls detail
Virtual Private LAN Service Instance: vpls2, ID: 2
 SIG-Protocol: LDP
  Route-Distinguisher :10.10.10.10:1001
  Route-Target :2:100
  L2 VPN ID :200:1001
 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: ce4/2.2
   Status: Up
   Subinterface Match Criteria(s) :
   dot1q 2

 Mesh Peers:
   160.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime:
00:02:55)
    FEC signaling element: FEC129
    FEC129 details:
      agi : 00 0A 00 C8 00 00 03 E9
      saii: 150.1.1.1
      taii: 160.1.1.1

   170.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime:
00:00:58)
    FEC signaling element: FEC129
    FEC129 details:
      agi : 00 0A 00 C8 00 00 03 E9
      saii: 150.1.1.1
      taii: 170.1.1.1

 Spoke Peers:
   vc1 (Up) (UpTime 00:03:28)


9600-PE3-Pri-HUB#sh bgp l2vpn vpls ldp-sig
VPLS-ID     Discovered-Peers  Route-Distinguisher    L2VPN-ID
2           2                 10.10.10.10:1001       200:1001
9600-PE3-Pri-HUB#sh bgp l2vpn vpls ldp-sig detail
VPLS-ID: 2
Local L2VPN-ID      : 200:1001
Local RD            : 10.10.10.10:1001
Local Route-Targets :
    Import List     : 2:100
    Export List     : 2:100
Discovered Peers    : 2
Mesh Peers          :
    BGP Peer-1      : 100.1.1.1
      Peer L2VPN-ID       : 200:1001
      Peer Route-Targets  :
        Export List       : 2:100
```

```
      Peer Up time        : 00:02:03

    BGP Peer-2      : 100.1.1.1
      Peer L2VPN-ID       : 200:1001
      Peer Route-Targets  :
        Export List       : 2:100
      Peer Up time        : 00:04:10


9600-PE3-Pri-HUB#
------------------------------------------------------------------------------------------
----------------------------------------

7535-6-PE4-Bkp-HUB#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
2          150.1.1.1         24325         26241      xe24           26241      2/Up
0          LDP           Active     00:02:33      -
2          170.1.1.1         25614         26242      xe24           26241      2/Up
0          LDP           Active     00:00:36      -
7535-6-PE4-Bkp-HUB#
7535-6-PE4-Bkp-HUB#
7535-6-PE4-Bkp-HUB#sh mpls vpls detail
Virtual Private LAN Service Instance: vpls2, ID: 2
 SIG-Protocol: LDP
  Route-Distinguisher :10.10.10.10:1001
  Route-Target :2:100
  L2 VPN ID :200:1001
 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: No) (Up) (UpTime:
00:03:00)
    FEC signaling element: FEC129
    FEC129 details:
      agi : 00 0A 00 C8 00 00 03 E9
      saii: 160.1.1.1
      taii: 150.1.1.1

   170.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime:
00:01:03)
    FEC signaling element: FEC129
    FEC129 details:
      agi : 00 0A 00 C8 00 00 03 E9
      saii: 160.1.1.1
      taii: 170.1.1.1

 Spoke Peers:
   vc1001 (Dn) (Reason: VC on standby)


7535-6-PE4-Bkp-HUB# sh bgp l2vpn vpls ldp-sig
VPLS-ID     Discovered-Peers  Route-Distinguisher    L2VPN-ID
2           2                 10.10.10.10:1001       200:1001
7535-6-PE4-Bkp-HUB#sh bgp l2vpn vpls ldp-sig detail
VPLS-ID: 2
Local L2VPN-ID      : 200:1001
Local RD            : 10.10.10.10:1001
Local Route-Targets :
    Import List     : 2:100
    Export List     : 2:100
Discovered Peers    : 2
Mesh Peers          :
    BGP Peer-1      : 100.1.1.1
      Peer L2VPN-ID       : 200:1001
      Peer Route-Targets  :
        Export List       : 2:100
      Peer Up time        : 00:02:14

    BGP Peer-2      : 100.1.1.1
      Peer L2VPN-ID       : 200:1001
      Peer Route-Targets  :
        Export List       : 2:100
      Peer Up time        : 00:04:25


-------------------------------------------------------------------------------------------
-----------------------------------------
7535-7-PE5#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
2          150.1.1.1         25618         26240      xe24           26242      2/Up
0          LDP           Active     00:00:08      -
2          160.1.1.1         25619         26241      xe24           26242      2/Up
0          LDP           Active     00:00:08      -
7535-7-PE5#
7535-7-PE5#
7535-7-PE5#
7535-7-PE5#sh mpls vpls detail
Virtual Private LAN Service Instance: vpls2, ID: 2
 SIG-Protocol: LDP
  Route-Distinguisher :10.10.10.10:1001
  Route-Target :2:100
  L2 VPN ID :200:1001
 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: No) (Up) (UpTime:
00:00:26)
    FEC signaling element: FEC129
    FEC129 details:
      agi : 00 0A 00 C8 00 00 03 E9
      saii: 170.1.1.1
      taii: 150.1.1.1

   160.1.1.1 (Type: Ethernet VLAN) (Negotiated - CW: No, FAT: No) (Up) (UpTime:
00:00:26)
    FEC signaling element: FEC129
    FEC129 details:
      agi : 00 0A 00 C8 00 00 03 E9
      saii: 170.1.1.1
      taii: 160.1.1.1
```

```
7535-7-PE5#sh bgp l2vpn vpls ldp-sig
VPLS-ID     Discovered-Peers  Route-Distinguisher    L2VPN-ID
2           2                 10.10.10.10:1001       200:1001
7535-7-PE5#sh bgp l2vpn vpls ldp-sig detail
VPLS-ID: 2
Local L2VPN-ID      : 200:1001
Local RD            : 10.10.10.10:1001
Local Route-Targets :
    Import List     : 2:100
    Export List     : 2:100
Discovered Peers    : 2
Mesh Peers          :
    BGP Peer-1      : 100.1.1.1
      Peer L2VPN-ID       : 200:1001
      Peer Route-Targets  :
        Export List       : 2:100
      Peer Up time        : 00:02:30

    BGP Peer-2      : 100.1.1.1
      Peer L2VPN-ID       : 200:1001
      Peer Route-Targets  :
        Export List       : 2:100
      Peer Up time        : 00:02:30

------------------------------------------------------------------------------------------
----------------------------------------
7946-74-PE2-SPOKE#sh 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: dot1.q
 Operating mode: Raw
 MAC Withdrawal:

 Configured interfaces:
  Interface: xe38.2
   Status: Up
   Subinterface Match Criteria(s) :
   dot1q 2

 Spoke Peers:
   vc1 (Up) (UpTime 00:04:18)
    Secondary: vc1001 (Dn) (Reason: VC on standby)
```

```
7946-74-PE2-SPOKE#sh run vpls
!
mpls vpls vpls2 2
 vpls-mtu 9000
 vpls-vc vc1
  secondary vc1001
  exit-spoke
 exit-vpls
!
!
interface xe38.2 switchport
 access-if-vpls
  mpls-vpls vpls2
!
7946-74-PE2-SPOKE#sh 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             0            26240      ce8/1          26240
2/Up     vc1001
2          160.1.1.1    vc1001          N/A          26241      N/A            26240
0/Dn     ---
7946-74-PE2-SPOKE#
```

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 6.0.0.

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 6.0.0.

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 both to the BGP AD VPLS.

Note: Only one route-target of type "both" is supported, and once it is configured, it cannot be modified. To change it, the entire bgp-auto-discovery configuration block must be removed.

Command Syntax

```
route-target (both) (ASN:nn_or_IP-address:nn|)
```

Parameters

`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 6.0.0.

Examples

```
#configure terminal
(config)#mpls vpls VPLS100 100
(config-vpls)#signaling ldp
(config-vpls-sig)#bgp-auto-discovery
(config-vpls-ldp-sig-bgp-ad)#route-target both 2.2.2.2:2222
(config-vpls-ldp-sig-bgp-ad)#exit-bgp-auto-discovery
(config-vpls-sig)#exit
(config-vpls)#exit
```

Improved Management

This section describes the network monitoring enhancements and new features introduced in the Release 6.6.0.

Enhanced Security and Performance

This section describes the security, performance, scalability, and access control enhancements and new features introduced in the Release 6.6.0.

- Y.1564 - Ethernet Service Activation Test Methodology

CHAPTER 1 Y.1564 - Ethernet Service Activation Test Methodology

Overview

This document describes ITU-T Y.1564, a standard for Ethernet Service Activation Testing (SAT). It focuses on validating Service Level Agreements (SLAs) through Key Performance Indicators (KPIs) such as Frame Delay (FD), Frame Loss Ratio (FLR), and Frame Delay Variation (FDV) in a single test. The document also covers traffic classification, service prioritization using Quality of Service (QoS), and performance verification mechanisms like Loopback Messages (LBM).

The SAT feature allows service providers to verify Ethernet service configurations and performance before activating them for customers.

OcNOS switches execute SAT by generating test traffic, collecting performance data, and ensuring SLA compliance before activation. This supports validation for both upstream (traffic entering the network interface) and downstream (traffic exiting the network interface) directions, with results displayed post-test.

The SAT allows service providers to verify Ethernet service configurations and performance before activating them for customers.

Currently, this feature is supported on OcNOS devices equipped with the Qumran2 and Jericho2 chipsets.

Feature Characteristics

An SLA (Service Level Agreement) is a binding agreement between a service provider and a customer, ensuring that the delivered service meets agreed-upon performance levels.

To meet SLA obligations, ITU-T Y.1564 provides a standardized procedure for verifying SLAs by:

Validating service configuration: Ensuring each Ethernet service is correctly set up.

Validating service quality: Ensuring the delivered quality meets user expectations.

Each User Network Interface (UNI) runs multiple services qualified by KPIs. The SAT verifies the network's ability to handle traffic within design parameters, ensuring compliance with Bandwidth Rate Profiles and Performance Criteria.

Bandwidth Profile

The SAT generates traffic based on the configured bandwidth profiles:

- *Committed Information Rate (CIR):* The guaranteed bandwidth available at all times for a specific service.
- *Excess Information Rate (EIR):* Additional bandwidth above the CIR, available based on network load.

OcNOS applies bandwidth profiles that include KPIs to indicate minimum performance requirements. Both CIR and EIR are essential for maintaining SLA compliance.

Performance Criteria

The SAT collects the following KPIs to ensure SLA requirements are met:

- *Frame Delay (FD):* Measures latency, the time taken for a packet to travel from source to destination.
- *Frame Delay Variation (FDV):* Also known as jitter, it measures variability in packet arrival times.
- *Frame Loss Ratio (FLR):* The percentage of packets lost due to errors or network congestion.

Color Codes for Traffic Policing

OcNOS uses traffic color codes to police ingress packets and manage their priority:

- Green: Traffic guaranteed at all times for a specific service.
- Yellow: Best effort traffic, utilizing excess bandwidth when available.
- Red: Traffic dropped without disrupting guaranteed services.

Quality of Service (QoS)

QoS ensures efficient traffic prioritization in the network. By assigning priority levels to each service and using appropriate prioritization algorithms, OcNOS delivers higher service quality for critical traffic. QoS differentiates traffic using specific fields in frames, ensuring higher-priority traffic receives preferential treatment.

Service Traffic Test Types

OcNOS performs the SAT on interfaces provisioned with rate profiles. Performance statistics are collected during the test (either locally at PE1 by looping the traffic back from PE2, if the test is two-way or at PE2, if the test is one-way.), which operates in either:

Upstream: Traffic is generated at the UNI and forwarded towards the network (From PE1 UNI to PE2 UNI).

Downstream: Traffic exit at the interface (PE1 network interface to PE2 network interface).

The illustration below depicts the traffic flow of Upstream and Downstream.

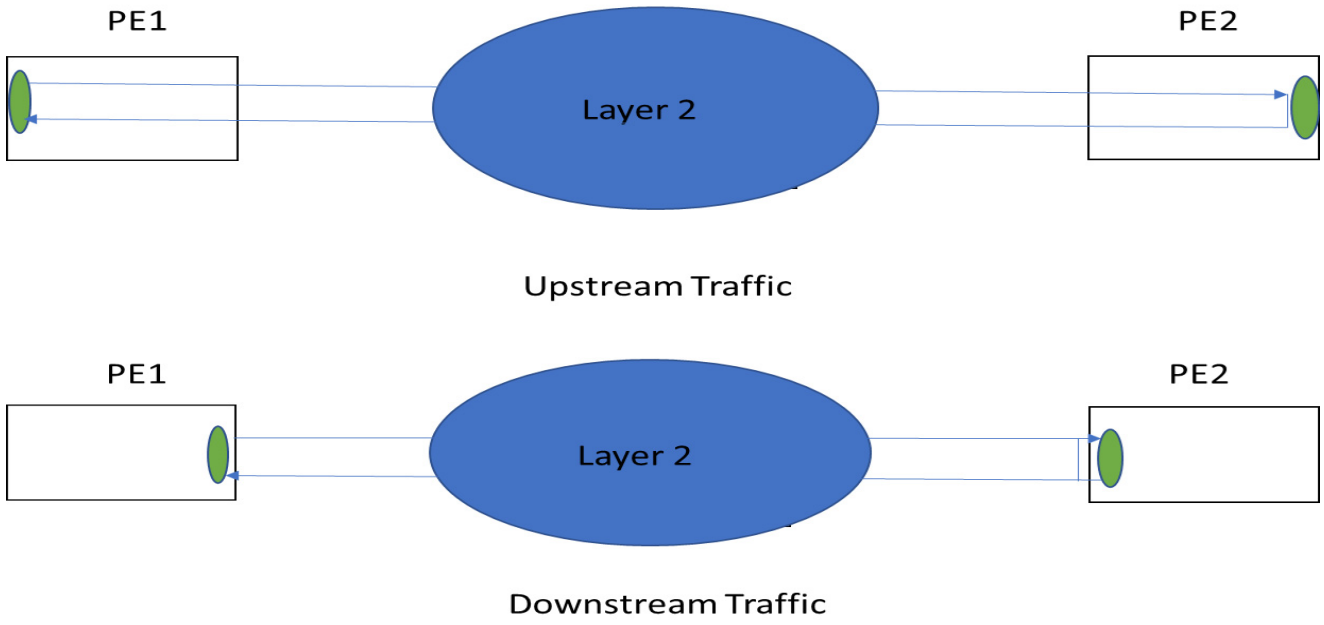

Figure 1-2: Upstream and Downstream Traffic Flow

SAT Performance Statistics

During SAT testing, the following statistics are collected to verify SLA compliance:

- Frame counts
- Errors
- Out-of-order frames

- Lost frames
- Minimum/maximum frame delay
- Frame delay variation

In two-way tests, Loopback Messages (LBM) and Loopback Replies (LBR) are used to measure throughput and frame loss. LBMs are sent as test traffic, while the peer device responds with LBR packets, allowing OcNOS to collect performance data.

SAT Profile Configuration

To perform SAT testing, a rate profile must be configured. It includes parameters such as CIR, EIR, Class of Service (CoS), and Drop Eligibility Indicator (DEI). Multiple test streams can be configured and run sequentially or in parallel.

For delay measurement, Delay Measurement Messages (DMMs) are used:

- When EIR is configured, VLAN priorities alternate between CIR CoS and EIR CoS.
- If EIR is not configured, only CIR CoS is used.

SAT profiles allow for the following operations:

- Create, abort, delete, or start a test
- Clear test history and results
- Configure frame size profiles for non-default frame sizes
- Accept or reject frame loss and delay criteria for CIR traffic
- Delete SAT rate profile or acceptance profile
- Create SAT rate profile or acceptance profile
- Delete SAT frame size profile
- Display SAT test summary and detailed SAT test results

Limitations:

- Supported only on Qumran2 and Jericho2C+ chipsets.
- Limited to 2-way traffic testing.
- CRC validation is not supported for SAT frames due to hardware limitations.
- LBM/LBR supports only a single TLV length, which may cause malformed frames for multiple frame sizes.
- LAG load balancing is supported only for UP MEPs; DOWN MEPs are unsupported.
- Supports a maximum of 8 CIR-EIR stream pairs, allowing up to 32 streams in various configurations. For example, you can configure 8 CIR-EIR pairs alongside 16 CIR-only streams, or alternatively, up to 32 CIR-only streams. This flexibility enables a variety of traffic configurations to accommodate different traffic scenarios.

Benefits

The ITU-T Y.1564-based SAT feature ensures Ethernet services meet SLA requirements before deployment. By validating configuration and performance, it minimizes service disruptions and reduces troubleshooting efforts. The ability to measure KPIs like latency, jitter, and frame loss enables precise SLA verification, resulting in:

- Higher customer satisfaction
- Improved network reliability
- Cost-effective service activation testing

Prerequisites

To implement ITU-T Y.1564 Service Activation Testing (SAT), the following prerequisites must be met:

- **Device Support** - Network devices (e.g., routers, switches) must support ITU-T Y.1564 standards and the necessary traffic generation and measurement features.
- **Configured Network Path** - A properly configured network path for test traffic between the source and destination endpoints is required.
- **Acceptance Criteria** - Clearly defined acceptance criteria for SLAs and performance KPIs (e.g., Frame Delay, Frame Loss Ratio, Frame Delay Variation) to validate test results.

Configuration

Following are the steps required to configure SAT:

1. Configuring rate profiles (CIR, EIR, QoS) that defines the minimum performance standards to ensure.
2. Configuring acceptance profiles (CIR, FLR, FD, FDV). The service must operate at or above the performance levels defined in the acceptance profile to be considered for bringing into service.
3. Configuring frame size profile.
4. Configure a SAT test.
5. Setting up measurement tools, such as Loopback Message (LBM), Loopback Reply (LBR), and Delay Measurement Message (DMM). The packets are generated and responded by OcNOS devices.
6. Executing SAT test by generating traffic in upstream/downstream directions.
7. Collecting and analyzing KPIs.

Topology

The topology uses a EVPN MPLS ELINE services between PE1 and PE2 nodes. It consists of two CE1 and CE2 nodes. Traffic that complies with rate profile is generated on PE1 and sent out through the interface where the service is provisioned.

SAT Topology for EVPN MPLS ELINE Service

Procedure to configure SAT

The procedure given below performs SAT for EVPN MPLS ELINE service. It assumes that the EVPN MPLS ELINE service exist between PE1 and PE2. Ensure the CFM up mep session exists between PE1 and PE2 nodes along with delay-measurement reply on the PE2. Refer to the *Y.1731 and CFM Over EVPN ELINE Single Home* chapter.

Configure PE1

1. Login to PE.

2. Set up SAT rate profile with rate parameters such as profile name, traffic direction two-way, mode generator-receiver, frame transfer delay times, color of the traffic frame is known, CIR value as 1000mbps, CoS 6 DEI 0, EIR value as 500 mbps, CoS 4 DEI 1 and the number of steps required to send the traffic load-step as 1. Following is the sample configuration.

```
(config)#ethernet sat rate-profile r1
(sat_rate_config_profile)#two-way
(sat_rate_config_profile)#mode generator-receiver
(sat_rate_config_profile)#delay-bins min us 1 max ms 1
(sat_rate_config_profile)#color-aware
(sat_rate_config_profile)#cir 1000 mbps cos 6 dei 0
(sat_rate_config_profile)#eir 500 mbps cos 4 dei 1
(sat_rate_config_profile)#load-step 1
(sat_rate_config_profile)#commit
```

3. Create SAT acceptance profile with the performance criteria such as FLR as 0.01, FLR policing as 25.00, FD as 100ms, FDV as 10ms. Following is the sample configuration.

```
(sat_rate_config_profile)#ethernet sat acceptance-profile a1
(sat_acceptance_config_profile)#cir frame-loss-ratio 0.0001
(sat_acceptance_config_profile)#cir frame-loss-ratio-policing 25.0000
(sat_acceptance_config_profile)#cir frame-delay-variation ms 10
(sat_acceptance_config_profile)#cir frame-delay ms 100
(sat_acceptance_config_profile)#commit
```

4. Set Frame size profile with various frame sizes for testing. Following is the sample configuration.

```
(sat_acceptance_config_profile)#ethernet sat frame-size-profile f1
(sat_frame_size_config_profile)#size-a 64
(sat_frame_size_config_profile)#size-b 128
(sat_frame_size_config_profile)#size-c 256
(sat_frame_size_config_profile)#size-d 512
(sat_frame_size_config_profile)#size-e 1024
(sat_frame_size_config_profile)#size-f 1280
(sat_frame_size_config_profile)#size-g 1518
(sat_frame_size_config_profile)#size-u 200
(sat_frame_size_config_profile)#commit
```

5. Configure SAT Test parameters such as traffic direction, streams type, service test for CIR-EIR, service stream, rate profile, acceptance profile, frame size profile, frame sequence, frame data pattern, frame payload, attach configured CFM MD, MA, MEP-ID, and destination MAC as learn from remote PE2 MAC. Following is the sample configuration.

```
(sat_frame_size_config_profile)#ethernet sat sat-test-1
(sat-test)#direction upstream
(sat-test)#stream-run parallel
(sat-test)#service-test cir-eir
(sat-test-service)#service-stream s1
(sat-test-stream)#rate-profile r1
(sat-test-stream)#acceptance-profile a1
(sat-test-stream)#frame-size-profile f1
(sat-test-stream)#frame-sequence bcd
(sat-test-stream)#frame-data-pattern null-sig
(sat-test-stream)#frame-payload
(sat-test-stream-payload)#cfm domain md000 ma ma10 mepid 10
```

```
(sat-test-stream-payload)#dst-mac e8c5.7adb.d59b
(sat-test-stream-payload)#commit
```

6. Execute the SAT test.

```
OcNOS#sat start sat-test-1
```

7. Validate and monitor the SAT test results. To view test summary, execute the below CLI.

```
OcNOS#show ethernet sat summary
```

8. To check the detailed SAT test results, execute the below CLI.

```
OcNOS#show ethernet sat detail sat-test-1
```

9. Abort or clear SAT tests, if required. To abort a running test, execute the below CLI.

```
OcNOS#sat abort sat-test-1
```

10. To clear SAT test data, execute the below CLI.

```
OcNOS#sat clear sat-test-1
```

11. Repeat test for failed streams. Following is the sample command to repeat the test.

```
OcNOS#sat start sat-test-1 repeat-last-fail-test
```

Validation

The following show output list the test summary of the SAT.

```
OcNOS#show ethernet sat summary
#Name Run Status
------------------------------------------------------------------------
sat-test-1 1 Passed (2019 Feb 14 10:16:23 - 2019 Feb 14 10:18:25)
sat-test-1 2 Aborted (2019 Feb 14 10:46:44 - 2019 Feb 14 10:48:04)
sat-test-1 3 Aborted (2019 Feb 14 10:52:13 - 2019 Feb 14 10:52:37)
sat-test-1 4 In-progress (2019 Feb 14 10:53:43 - )
OcNOS#
```

The following show output lists the detail report of the SAT.

```
OcNOS#! Show SAT test detail results
OcNOS#show ethernet sat detail sat-test-1
------------------------------------------------------------------------
SAT sat-test-1, Run 4
------------------------------------------------------------------------
Description : (Not Specified)
Oper state : Passed
Start Time : 2019 Feb 14 10:53:43
End Time : 2019 Feb 14 10:55:46
Direction : upstream

Stream Run : parallel
Stream s1 Config
------------------------------------------------------------------------
Bandwidth profile : CIR 1000 mbps, COS 6, DEI 0
EIR 500 mbps, COS 4, DEI 1
Color aware : on, color-method (DEI)
acceptance : CIR - FLR NA FTD 100 ms, FDV 10 ms
```

```
MTU : 128, 256, 512,
CFM : domain md000, ma ma10, mep 10
Destination e8c5.7adb.d59b
Stream s1, Test (cir-eir)
-------------------------------------------------
Oper State        : Failed
Test Duration     : 00:02:03 (hh:mm:ss)
Start Time        : 2019 Feb 14 10:53:43
End Time          : 2019 Feb 14 10:55:46
CIR:
Tx packets 50720289, bytes 15249900226
Rx packets 25445881, bytes 8694009171
FL 25274408, FLR 49.8309 %
FD 11 us, FDV 1 us
FD Min 10 us, FD Max 34 us, FD exceeded 0%
Out of Order packets 0
Error packets 0
EIR:
Tx packets 25358536, bytes 7624466320
Rx packets 12722142, bytes 4346731850
FL 12636394, FLR 49.8309 %
FD 11 us, FDV 1 us
FD Min 10 us, FD Max 36 us, FD exceeded 0%
Out of Order packets 0
Error packets 0
OcNOS#
```

The following shows the output of the SAT test executed again.

```
OcNOS# ! Run SAT test again
OcNOS#sat start sat-test-1
Name Run Status
-------------------------------------------------------------------------
sat-test-1 1 Passed (2019 Feb 14 10:16:23 - 2019 Feb 14 10:18:25)
sat-test-1 2 Aborted (2019 Feb 14 10:46:44 - 2019 Feb 14 10:48:04)
sat-test-1 3 Aborted (2019 Feb 14 10:52:13 - 2019 Feb 14 10:52:37)
sat-test-1 4 Passed (2019 Feb 14 10:53:43 - 2019 Feb 14 10:55:46)
OcNOS#
OcNOS#

OcNOS# ! Start and abort SAT test
OcNOS#sat start sat-test-1
OcNOS#sat abort sat-test-1
OcNOS#
OcNOS#! Show SAT test summary
OcNOS#show ethernet sat summary
OcNOS#
OcNOS#show ethernet sat summary
Name Run Status
-------------------------------------------------------------------------
sat-test-1 1 Passed (2019 Feb 14 10:16:23 - 2019 Feb 14 10:18:25)
```

```
sat-test-1 2 Aborted (2019 Feb 14 10:46:44 - 2019 Feb 14 10:48:04)
sat-test-1 3 Aborted (2019 Feb 14 10:52:13 - 2019 Feb 14 10:52:37)
sat-test-1 4 Passed (2019 Feb 14 10:53:43 - 2019 Feb 14 10:55:46)
sat-test-1 5 Aborted (2019 Feb 14 11:01:12 - 2019 Feb 14 11:01:54)
OcNOS#
OcNOS#

OcNOS#! Show SAT test detail results
OcNOS#show ethernet sat detail sat-test-1
------------------------------------------------------------------------
SAT sat-test-1, Run 5
------------------------------------------------------------------------
Description : (Not Specified)
Oper state : Passed
Start Time : 2019 Feb 14 11:01:12
End Time : 2019 Feb 14 11:01:54
Direction : upstream
Stream Run : parallel
Stream s1 Config
------------------------------------------------------------------------
Bandwidth profile : CIR 1000 mbps, COS 6, DEI 0
EIR 500 mbps, COS 4, DEI 1
Color aware : on, color-method (DEI)
acceptance : CIR - FLR NA FTD 100 ms, FDV 10 ms
MTU : 128, 256, 512,
CFM : domain md000, ma ma10, mep 10
Destination e8c5.7adb.d59b
Stream s1, Test (cir-eir)
-----------------------------------------------------
Oper State : Aborted
Test Duration : 00:00:42 (hh:mm:ss)
Start Time : 2019 Feb 14 11:01:12
End Time : 2019 Feb 14 11:01:54
CIR:
Tx packets 16716322, bytes 5026040644
Rx packets 16716322, bytes 5711409846
FL 0, FLR 0.0000
FD 12 us, FDV 2 us
FD Min 10 us, FD Max 36 us, FD exceeded 0%
Out of Order packets 0
Error packets 0
EIR:
Tx packets 8357666, bytes 2512871364
Rx packets 8357666, bytes 2855535670
FL 0, FLR 0.0000
FD 0 us, FDV 0 us
FD Min 0 us, FD Max 0 us, FD exceeded 0%
Out of Order packets 0
Error packets 0
OcNOS#
```

Implementation Examples

Example 1

The following is a sample SAT configuration for CIR 38000mbps with the frame size of 1518 bytes for the duration of 1 minute.

```
ethernet sat rate-profile rate6
  delay-bins min us 1 max ms 1
  cir 380000 mbps cos 0 dei 0
!
ethernet sat acceptance-profile profile1
  cir frame-loss-ratio 0.0001
  cir frame-loss-ratio-policing 25.0000
  cir frame-delay-variation ms 10
  cir frame-delay ms 10
!
ethernet sat frame-size-profile size1
  size-a 64
  size-b 128
  size-c 256
  size-d 512
  size-e 1024
  size-f 1280
  size-g 1518
  size-h 9200
  size-u 80
!
ethernet sat test1
  direction upstream
  stream-run one-by-one
  service-test performance
    duration minutes 1
  service-stream stream1
    rate-profile rate6
    acceptance-profile profile1
    frame-size-profile size1
    frame-sequence g
    frame-payload
      cfm domain md001 ma ma001 mepid 444
      dst-mac  e001.a6b8.ed05
!
```

The following shows the output of the above SAT.

```
CE2#show ethernet sat detail test1
------------------------------------------------------------------------
SAT test1, Run 1
------------------------------------------------------------------------
Description      : (Not Specified)
```

```
Oper state       : Passed
Start Time       : 2025 Jan 06 13:41:14
End Time         : 2025 Jan 06 13:42:18
Direction        : upstream
Stream Run       : one-by-one

Stream stream1 Config
------------------------------------------------------------------------
Bandwidth profile : CIR 380000 mbps, COS 0, DEI 0
                    EIR 0 mbps, COS 0, DEI 0
Color aware       : off, color-method (DEI)
Acceptance        : CIR - FLR 0.0001 %, FTD 10 ms, FDV 10 ms
MTU               : 1518,
Frame Info        : FD enabled, prbs pattern
CFM               : domain md001, ma ma001, mep 444
                    Destination e001.a6b8.ed05

Stream stream1, Test (performance)
-----------------------------------------------------
Oper State        : Passed
Test Duration     : 00:01:00 (hh:mm:ss)
Start Time        : 2025 Jan 06 13:41:14
End Time          : 2025 Jan 06 13:42:14

CIR:
Tx packets 1869134626, bytes 2837346362268
Rx packets 1869134626, bytes 2837346362268
FL 0, FLR 0.0000 %
FD 27 us, FDV 0 us
FD Min 25 us, FD Max 31 us, FD exceeded 0%
Out of Order packets 0
Error packets 1

EIR:
Tx packets 0, bytes 0
Rx packets 0, bytes 0
FL 0, FLR 0.0000 %
FD 19 us, FDV 0 us
FD Min 18 us, FD Max 19 us, FD exceeded 0%
Out of Order packets 0
Error packets 0
```

Example 2

The following is a sample SAT configuration for CIR and EIR 1000mbps with the frame delay of 1 ms.

```
PE1-7038-Ufi-Q2C#show ru sat
!
ethernet sat rate-profile rate1
  delay-bins min ms 1 max ms 1
```

```
  color-aware
  cir 1000 mbps cos 4 dei 0
  eir 1000 mbps cos 5 dei 1
!
ethernet sat acceptance-profile profile1
  cir frame-loss-ratio 0.0001
  cir frame-loss-ratio-policing 25.0000
  cir frame-delay-variation ms 10
  cir frame-delay ms 10
!
ethernet sat frame-size-profile size1
  size-a 64
  size-b 128
  size-c 256
  size-d 512
  size-e 1024
  size-f 1280
  size-g 1518
  size-h 9200
  size-u 80
!
ethernet sat test1
  direction upstream
  stream-run parallel
  service-test cir
    duration minutes 1
  service-test cir-eir
    duration minutes 1
  service-stream stream1
    rate-profile rate1
    acceptance-profile profile1
    frame-size-profile size1
    frame-sequence bcd
    frame-payload
      cfm domain test1 ma ma100 mepid 1001
      dst-mac 5c07.5851.c9d5
!

PE1-7038-Ufi-Q2C#show ethernet sat summary
Name                         Run  Status
-------------------------------------------------------------------------------------
test1                       1    Aborted     (2025 Jan 17 14:44:26 - 2025 Jan 17 14:44:39)
test1                       2    Passed      (2025 Jan 17 14:44:50 - 2025 Jan 17 14:46:57)
test1                       3    Passed      (2025 Jan 17 14:49:51 - 2025 Jan 17 14:50:54)
PE1-7038-Ufi-Q2C#

PE1-7038-Ufi-Q2C#show ethernet sat detail test1
-----------------------------------------------------------------------
SAT test1, Run 2
-----------------------------------------------------------------------
```

```
Description      : (Not Specified)
Oper state       : Passed
Start Time       : 2025 Jan 17 14:44:50
End Time         : 2025 Jan 17 14:46:57
Direction        : upstream
Stream Run       : parallel

Stream stream1 Config
------------------------------------------------------------------------
Bandwidth profile : CIR 1000 mbps, COS 4, DEI 0
                    EIR 1000 mbps, COS 5, DEI 1
Color aware       : on, color-method (DEI)
Acceptance        : CIR - FLR 0.0001 %, FTD 10 ms, FDV 10 ms
MTU               : 128, 256, 512,
Frame Info        : FD enabled, prbs pattern
CFM               : domain test1, ma ma100, mep 1001
                    Destination 5c07.5851.c9d5

Stream stream1, Test (cir-step-1) Load 100%
-----------------------------------------------------
Oper State        : Passed
Test Duration     : 00:01:00 (hh:mm:ss)
Start Time        : 2025 Jan 17 14:44:50
End Time          : 2025 Jan 17 14:45:50

Tx packets 24945535, bytes 7450399616
Rx packets 24945535, bytes 7450399616
FL 0, FLR 0.0000 %
FD 12 us, FDV 2 us
FD Min 11 us, FD Max 41 us, FD exceeded 0%
Out of Order packets 0
Error packets 0

Stream stream1, Test (cir-eir)
-----------------------------------------------------
Oper State        : Passed
Test Duration     : 00:01:00 (hh:mm:ss)
Start Time        : 2025 Jan 17 14:45:54
End Time          : 2025 Jan 17 14:46:54

CIR:
Tx packets 24948497, bytes 7451284224
Rx packets 24948497, bytes 7451284224
FL 0, FLR 0.0000 %
FD 16 us, FDV 3 us
FD Min 11 us, FD Max 52 us, FD exceeded 0%
Out of Order packets 0
Error packets 0

EIR:
```

```
Tx packets 24948497, bytes 7451284224
Rx packets 24948497, bytes 7451284224
FL 0, FLR 0.0000 %
FD 11 us, FDV 0 us
FD Min 11 us, FD Max 11 us, FD exceeded 0%
Out of Order packets 0
Error packets 0

PE1-7038-Ufi-Q2C#sat start test1 cir-eir service-stream stream1
PE1-7038-Ufi-Q2C#2025 Jan 17 14:50:55.784 : PE1-7038-Ufi-Q2C : ONMD : INFO : [ETH SAT
STATUS CHANGE_5]: Tests completed for the test1 and run is 3

PE1-7038-Ufi-Q2C#show ethernet sat detail test1
------------------------------------------------------------------------
SAT test1, Run 3
------------------------------------------------------------------------
Description      : (Not Specified)
Oper state       : Passed
Start Time       : 2025 Jan 17 14:49:51
End Time         : 2025 Jan 17 14:50:54
Direction        : upstream
Stream Run       : parallel

Stream stream1 Config
------------------------------------------------------------------------
Bandwidth profile : CIR 1000 mbps, COS 4, DEI 0
                    EIR 1000 mbps, COS 5, DEI 1
Color aware       : on, color-method (DEI)
Acceptance        : CIR - FLR 0.0001 %, FTD 10 ms, FDV 10 ms
MTU               : 128, 256, 512,
Frame Info        : FD enabled, prbs pattern
CFM               : domain test1, ma ma100, mep 1001
                    Destination 5c07.5851.c9d5

Stream stream1, Test (cir-eir)
------------------------------------------------
Oper State        : Passed
Test Duration     : 00:01:00 (hh:mm:ss)
Start Time        : 2025 Jan 17 14:49:51
End Time          : 2025 Jan 17 14:50:51

CIR:
Tx packets 24945587, bytes 7450415104
Rx packets 24945587, bytes 7450415104
FL 0, FLR 0.0000 %
FD 16 us, FDV 4 us
FD Min 11 us, FD Max 55 us, FD exceeded 0%
Out of Order packets 0
Error packets 0

EIR:
```

```
Tx packets 24945583, bytes 7450413952
Rx packets 24945583, bytes 7450413952
FL 0, FLR 0.0000 %
FD 11 us, FDV 0 us
FD Min 11 us, FD Max 11 us, FD exceeded 0%
Out of Order packets 0
Error packets 0


PE1-7038-Ufi-Q2C#
```

CLI Commands

The following configuration commands are introduced for SAT functionality.

- ethernet sat profile
- two-way
- mode generator-receiver | generator | receiver
- cir value cos dei
- eir VALUE cos dei
- color-aware
- delay-bins min max
- load-step
- sat acceptance-profile NAME
- cir frame-loss-ratio
- cir frame-loss-ratio-policing
- cir frame-delay-variation
- ethernet sat frame-size-profile NAME
- size-a|b|c|d|e|f|g|h|u
- ethernet sat NAME
- description WORD
- direction downstream | upstream
- stream-run
- service-test
- duration hours minutes seconds
- service-stream NAME
- rate-profile NAME
- acceptance-profile NAME
- frame-size-profile NAME
- frame-sequence WORD
- frame-data-pattern null-sig | prbs

- frame-payload
- cfm domain NAME ma NAME mepid <1-8191>
- dst-mac XXXX.XXXX.XXXX
- sat start NAME
- abort sat NAME
- clear sat NAME
- show ethernet sat summary
- show ethernet sat detail NAME

ethernet sat profile

Use this command to create Service Traffic Test (SAT) rate profile. A SAT rate profile defines various parameters such as traffic direction (one-way or two-way, refer to Service Traffic Test Types for more information.), test mode, CIR and/or EIR, and so on.

Note: In Jericho2 devices, a maximum of 8 CIR and EIR rate profiles and 32 CIR-only rate profiles are supported.

Command Syntax

```
ethernet sat profile NAME
```

Parameters

`sat profile NAME`

Specifies the name of the SAT rate profile.

Default

None

Command Mode

Configuration mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to create a SAT rate profile.

```
OcNOS(config)#ethernet sat rate-profile r1
```

two-way

Use this command to indicate that the traffic direction type is two-way.

Command Syntax

```
two-way
```

Parameters

None

Default

None

Command Mode

SAT_RATE_ CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the traffic direction is two-way.

```
OcNOS(sat-rate-config-profile)#two-way
```

mode generator-receiver | generator | receiver

Use this command to mention the traffic type. For a two-way test, use both the `generator-receiver`, whereas a one-way test requires only the `generator` or the `receiver`.

Command Syntax

```
mode generator-receiver | generator | receiver
```

Parameters

`mode generator-receiver`

Specifies the traffic type is both generator and receiver.

`mode generator`

Specifies the traffic type is generator only.

`mode receiver`

Specifies the traffic type is receiver only.

Default

None

Command Mode

SAT_RATE_ CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the test traffic is both generator and receiver.

```
OcNOS(sat-rate-config-profile)# mode generator-receiver
```

cir value cos dei

Use this command to configure the Committed Information Rate (CIR), Class of Service (COS) and Drop Eligibility Indicator (DEI).

Command Syntax

```
cir VALUE kbps|mbps cos <0-7>  dei <0-1>
```

Parameters

`cir VALUE`	Specifies the CIR value in kbps or mbps
`cos <0-7>`	Specifies the COS value for the class of traffic
`dei <0-1>`	Specify the DEI value ‘0’ not to drop the traffic and ‘1’ to drop the traffic.

Default

None

Command Mode

SAT_RATE_ CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the CIR, COS and DEI rate value in SAT rate profile.

```
OcNOS(sat-rate-config-profile)#cir 1000 mbps cos 6 dei 0
```

eir VALUE cos dei

Use this command to configure the Excess Information Rate (EIR), Class of Service (COS) and Drop Eligibility Indicator (DEI).

Command Syntax

```
eir VALUE kbps|mbps cos <0-7> dei <0-1>
```

Parameters

`eir VALUE`	Specifies the EIR value in kbps or mbps
`cos <0-7>`	Specifies the COS value for the class of traffic
`dei <0-1>`	Specify the DEI value ‘0’ not to drop the traffic and ‘1’ to drop the traffic.

Default

None

Command Mode

SAT_RATE_ CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the EIR, COS and DEI rate value in SAT rate profile.

```
OcNOS(sat-rate-config-profile)#eir 500 mbps cos 4 dei 1
```

color-aware

Use this command to configure the color of the traffic frame is known.

- Green Traffic: Guaranteed at all times for a specific service.
- Yellow Traffic: Best effort, utilized when excess bandwidth is available.
- Red Traffic: Dropped without disrupting services.

Command Syntax

```
color-aware
```

Parameters

None

Default

None

Command Mode

SAT_RATE_ CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure that the traffic bandwidth color is known.

```
OcNOS(sat-rate-config-profile)#color-aware
```

delay-bins min max

Use this command to define the minimum and maximum frame transfer delay times. This measures the time delay between a packet's transmission and its reception.

Command Syntax

```
delay-bins min ns|us|ms <1-4000> max ns|us|ms <1-4000>
```

Parameters

`delay-bins min ns|us|ms <1-4000>`

Specifies the minimum frame transfer delay time.

`max ns|us|ms <1-4000>`

Specifies the maximum frame transfer delay time.

Default

None

Command Mode

SAT_RATE_ CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the frame transfer delay time for the SAT rate profile.

```
OcNOS(sat-rate-config-profile)#delay-bins min us 1 max ms 1
```

load-step

Use this command to configure the load steps which defines the number of iterations required to send the traffic load.

For instance, if you want to send 100 Mbps of traffic and configure the load-step as 4, the traffic will be incremented in four stages until the full load is transferred.

The traffic is divided into four parts, starting with 25 Mbps (100 Mbps ÷ 4). Initially, 25 Mbps is sent. After completing the first test, the load increases to 50 Mbps for the second step. Once the second test finishes, the traffic is increased to 75 Mbps. Finally, the full load of 100 Mbps is sent in the last step.

Command Syntax

```
load-step VALUE
```

Parameters

```
load-step VALUE
```

Specifies the number of iterations required to send the traffic load.

Default

None

Command Mode

SAT_RATE_ CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the load-step value for SAT rate profile.

```
OcNOS(sat-rate-config-profile)#load-step 1
```

sat acceptance-profile NAME

Use this command to create an acceptance SAT profile.

Command Syntax

```
sat acceptance-profile NAME
```

Parameters

`acceptance-profile NAME`

Specifies the SAT acceptance profile name.

Default

None

Command Mode

SAT_RATE_CONFIG_PROFILE

SAT_ACCEPTANCE_ CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

The following command create a SAT acceptance profile.

```
OcNOS(sat-rate-config-profile)#ethernet sat acceptance-profile a1
OcNOS(sat-acceptance-config-profile)#
```

cir frame-loss-ratio

Use this command to configure the allowed frame loss ratio (FLR) for CIR traffic. The FLR refers to frames lost due to transmission errors or network congestion compared to the total frames sent.

Command Syntax

```
cir frame-loss-ratio <0.0001-100.0000>
```

Parameters

`cir frame-loss-ratio <0.0001-100.0000>`

Specifies the frame loss ratio value.

Default

None

Command Mode

SAT_ACCEPTANCE_CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the FLR value for CIR traffic in acceptance SAT profile.

```
OcNOS(sat-acceptance-config-profile)#cir frame-loss-ratio 0.01
```

cir frame-loss-ratio-policing

Use this command to policing the CIR frame lose ratio.

Command Syntax

```
cir frame-loss-ratio-policing <0.0001-100.0000>
```

Parameters

`cir frame-loss-ratio-policing <0.0001-100.0000>`

Specifies the frame loss ratio policing value

Default

None

Command Mode

SAT_ACCEPTANCE_CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to policing the FLR for CIR traffic in acceptance SAT profile.

```
OcNOS(sat-acceptance-config-profile)#cir frame-loss-ratio-policing 0.01
```

cir frame-delay

Use this command to configure allowed frame delay for CIR traffic.

Command Syntax

```
cir frame-delay ns|us|ms <1-10000>
```

Parameters

`cir frame-delay ns|us|ms <1-10000>`

Specifies the allowed CIR frame delay.

Default

None

Command Mode

SAT_ACCEPTANCE_CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to specify the frame delay value for CIR traffic in acceptance SAT profile.

```
OcNOS(sat-acceptance-config-profile)#cir frame-delay ms 100
```

cir frame-delay-variation

Use this command to configure the allowed frame delay variation value for CIR traffic.

The frame delay variation refers to the variability in arrival time between packet deliveries.

Command Syntax

```
cir frame-delay-variation ns|us|ms <1-10000>
```

Parameters

`cir frame-delay-variation ns|us|ms <1-10000>`

Specifies the allowed CIR frame delay variation.

Default

None

Command Mode

SAT_ACCEPTANCE_CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the accepted CIR frame delay variation value.

```
OcNOS(sat-acceptance-config-profile)#cir frame-delay-variation ms 10
```

ethernet sat frame-size-profile NAME

Use this command to create SAT frame size profile to change the SAT frame size to non default values.

Command Syntax

```
ethernet sat frame-size-profile NAME
```

Parameters

`frame-size-profile NAME`

Specifies the SAT frame size profile name.

Default

None

Command Mode

SAT_FRAME_SIZE_CONFIG_PROFILE

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to create a SAT frame size profile.

```
OcNOS(sat-acceptance-config-profile#ethernet sat frame-size-profile f1
OcNOS(sat-frame-size-config-profile#
```

size-a|b|c|d|e|f|g|h|u

Use this command to configure the frame size value to be used in the service and performance SAT test.

Command Syntax

```
size-a|b|c|d|e|f|g|h|u VALUE
```

Parameters

`size-a|b|c|d|e|f|g|h|u VALUE`

Specify frame size value using `size-a, size-b, size-c` etc. The following denotes the value of each size.

a	b	c	d	e	f	g	h	u
64	128	256	512	1024	1280	1518	MTU	User defined

Default

None

Command Mode

SAT FRAME SIZE CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the frame size to use in SAT test.

```
OcNOS(sat-acceptance-config-profile#ethernet sat frame-size-profile f1
OcNOS(sat-frame-size-config-profile#size-a 64
OcNOS(sat-frame-size-config-profile#size-b 128
OcNOS(sat-frame-size-config-profile#size-c 256
OcNOS(sat-frame-size-config-profile#size-d 512
OcNOS(sat-frame-size-config-profile#size-e 1024
OcNOS(sat-frame-size-config-profile#size-f 1280
OcNOS(sat-frame-size-config-profile#size-g 1518
```

ethernet sat NAME

Use this command to configure a name to SAT test.

Command Syntax

```
ethernet sat NAME
```

Parameters

`sat NAME`	Specify the name of SAT test.

Default

None

Command Mode

SAT TEST CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command configures a name to SAT test.

```
OcNOS(sat-frame-size-config-profile#ethernet sat sat-test-1
OcNOS(sat-test)#
```

description WORD

Use this command to provide a description to the SAT test.

Command Syntax

```
description WORD
```

Parameters

`description NAME`	Describes the SAT test.

Default

None

Command Mode

SAT_TEST_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command tp add a description to the SAT test.

```
OcNOS(sat-test)description sat-test for service activation
```

direction downstream | upstream

Use this command to configure the SAT test direction. Refer to Service Traffic Test Types section for more information.

Command Syntax

```
direction downstream | upstream
```

Parameters

`direction downstream`

Describes the traffic test direction is downstream.

`direction upstream`

Describes the traffic test direction is upstream.

Default

None

Command Mode

SAT_TEST_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure traffic test direction as upstream for SAT test.

```
OcNOS(sat-test)#direction upstream
```

stream-run

Use this command to configure SAT test streams type sequential or concurrent.

Command Syntax

```
stream-run one-by-one | parallel
```

Parameters

`stream-run one-by-one`

Describes the traffic stream is in sequential mode

`stream-run parallel`

Describes the traffic stream is in concurrent mode.

Default

None

Command Mode

SAT_TEST_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the traffic stream type as parallel.

```
OcNOS(sat-test)#stream-run parallel
```

service-test

Use this command to specify the services to use in SAT.

Command Syntax

```
service-test cir | cir-eir | performance | policing
```

Parameters

`service-test cir`

Specifies the CIR service test.

`service-test cir-eir`

Specifies the CIR-EIR service test.

`service-test performance`

Specifies the performance test.

`service-test policing`

Specifies the SAT policing test.

Default

None

Command Mode

SAT_TEST_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to include CIR-EIR service in SAT.

```
OcNOS(sat-test)#service-test cir-eir
```

duration hours minutes seconds

Use this command to configure SAT service test duration.

Command Syntax

```
duration hours <1-60> minutes <1-60> seconds <1-60>
```

Parameters

`hours <1-60>`	Specifies the service test duration in hours.
`minutes <1-60>`	Specifies the service test duration in minutes.
`seconds <1-60>`	Specifies the service test duration in seconds.

Default

None

Command Mode

SAT_TEST_SERVICE_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure a duration to execute the SAT.

```
OcNOS(sat-test-service)# duration minutes 2 seconds 2
```

service-stream NAME

Use this command to specify the service stream name to use in SAT.

Command Syntax

```
service-stream NAME
```

Parameters

`service-stream NAME`	Describes the service stream name.

Default

None

Command Mode

SAT_TEST_SERVICE_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to define the service stream name to use in SAT.

```
OcNOS(sat-test-service)#service-stream s1
```

rate-profile NAME

Use this command to specify the rate profile name to use in SAT.

Command Syntax

```
rate-profile NAME
```

Parameters

`rate-profile NAME`

Specifies the rate profile name to be used in the service test.

Default

None

Command Mode

SAT_TEST_STREAM_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the rate profile name to use in SAT.

```
OcNOS(sat-test-stream)#rate-profile r1
```

acceptance-profile NAME

Use this command to specify the name of the acceptance profile to use in SAT.

Command Syntax

```
acceptance-profile NAME
```

Parameters

`acceptance-profile NAME`

Specifies the acceptance profile name.

Default

None

Command Mode

SAT_TEST_STREAM_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to specify the acceptance profile name to use in SAT.

```
OcNOS(sat-test-stream)#acceptance-profile a1
```

frame-size-profile NAME

Use this command to specify the name of the frame size profile to use in SAT.

Command Syntax

```
frame-size-profile NAME
```

Parameters

`frame-size-profile NAME`

Specifies the frame size profile name.

Default

None

Command Mode

SAT_TEST_STREAM_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to specify the frame size profile name to use in SAT.

```
OcNOS(sat-test-stream)#frame-size-profile f1
```

frame-sequence WORD

Use this command to configure the frame sequence (`size-a|b|c|d|e|f|g|h|u`) to use in SAT.

Command Syntax

```
frame-sequence WORD
```

Parameters

`frame-sequence WORD`

Specifies the frame size sequence.

Default

None

Command Mode

SAT_TEST_STREAM_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to specify the frame size sequence to use in SAT.

```
OcNOS(sat-test-stream)#frame-sequence bcd
```

frame-data-pattern null-sig | prbs

Use this command to configure different frame size patterns for testing. The pattern defines a repeated sequence of frame sizes. For instance, the pattern "aabc" corresponds to frame sizes 64, 64, 128, and 256 (as an example for a data pattern).

Command Syntax

```
frame-data-pattern null-sig | prbs |
```

Parameters

`frame-data-pattern null-sig`

Specifies the frame data pattern sequence as null signal.

`frame-data-pattern prbs`

Specifies the frame data pattern sequence as pseudo random binary sequence.

Default

None

Command Mode

SAT_TEST_STREAM_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to specify frame pattern to use in SAT.

```
OcNOS(sat-test-stream)#frame-data-pattern null-sig
```

frame-payload

Use this command to configure the frame payload to use in SAT.

Command Syntax

```
frame-payload
```

Parameters

`frame-payload`

Specifies the frame payload.

Default

None

Command Mode

SAT_TEST_STREAM_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the frame payload to use in SAT.

```
OcNOS(sat-test-stream)#frame-payload
```

cfm domain NAME ma NAME mepid <1-8191>

Use this command to configure CFM domain, `ma` and `mepid` to use in SAT.

Command Syntax

```
cfm domain NAME ma NAME mepid <1-8191>
```

Parameters

`cfm domain NAME` Specifies the CFM domain name.

`ma NAME` Specifies the `ma` name.

`mepid <a-8191>` Specifies the `mepid` identifier.

Default

None

Command Mode

SAT_TEST_STREAM_PAYLOAD_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configure the CFM domain, ma and mepid to use in SAT.

```
OcNOS(sat-test-stream-payload)#cfm domain md000 ma ma10 mepid 10
```

dst-mac XXXX.XXXX.XXXX

Use this command to configure destination MAC address to use in SAT.

Command Syntax

```
dst-mac XXXX.XXXX.XXXX
```

Parameters

`dst-mac XXXX.XXXX.XXXX`

Specifies the CFM domain name.

Default

None

Command Mode

SAT_TEST_STREAM_PAYLOAD_CONFIG

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to configures the destination MAC address to use in SAT.

```
OcNOS(sat-test-stream-payload)#dst-mac e8c5.7adb.d59b
```

sat start NAME

Use this command to start a service activation test.

Command Syntax

```
sat start NAME <service test>
```

Parameters

`sat start NAME`

Specifies the service test name to start test.

`cir`	Executes cir test
`cir-eir`	Executes cir-eir test
`performance`	Executes performance test
`policing`	Executes policing test
`repeat-last-fail-test`	Repeat from the last failed SAT test case.
`service-stream`	Executes service stream test.

Default

None

Command Mode

Execute mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to start the SAT.

```
OcNOS#sat start sat-test-1
OcNOS#sat start sat-test-1 repeat-last-fail-test
```

abort sat NAME

Use this command to abort a service activation test.

Command Syntax

```
abort sat NAME
```

Parameters

`abort sat NAME`

Specifies the service test name to abort test.

Default

None

Command Mode

Execute mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to abort the running SAT.

```
OcNOS#sat abort sat-test-1
```

clear sat NAME

Use this command to clear the results of a service activation test.

Command Syntax

```
clear sat NAME
```

Parameters

`clear sat NAME`

Specifies the service test name to clear test results.

Default

None

Command Mode

Execute mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to clear the SAT results.

```
OcNOS#sat clear sat-test-1
```

show ethernet sat summary

Use this command to display the SAT results and status.

Command Syntax

```
show ethernet sat summary
```

Parameters

None

Default

None

Command Mode

Execute mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to display the SAT output.OcNOS#! Show SAT test summary

```
OcNOS#show ethernet sat summary
#Name                     Run  Status
-------------------------------------------------------------------------------------------
sat-test-1                1    Passed      (2019 Feb 14 10:16:23 - 2019 Feb 14 10:18:25)
sat-test-1                2    Aborted     (2019 Feb 14 10:46:44 - 2019 Feb 14 10:48:04)
sat-test-1                3    Aborted     (2019 Feb 14 10:52:13 - 2019 Feb 14 10:52:37)
sat-test-1                4    In-progress (2019 Feb 14 10:53:43 - )
OcNOS#
```

show ethernet sat detail NAME

Use this command to display specific SAT results detail.

Command Syntax

```
show ethernet sat detail NAME
```

Parameters

```
sat detail NAME
```

Specifies the service test name to show the detail results.

Default

None

Command Mode

Execute mode

Applicability

Introduced in OcNOS version 6.6.0.

Example

Command to display the SAT results detail.

```
OcNOS#! Show SAT test detail results
```

```
OcNOS#show ethernet sat detail sat-test-1
------------------------------------------------------------------------
SAT sat-test-1, Run 4
------------------------------------------------------------------------
Description      : (Not Specified)
Oper state       : Passed
Start Time       : 2019 Feb 14 10:53:43
End Time         : 2019 Feb 14 10:55:46
Direction        : upstream
Stream Run       : parallel
Stream s1 Config
------------------------------------------------------------------------
Bandwidth profile : CIR 1000 mbps, COS 6, DEI 0
                    EIR 500 mbps, COS 4, DEI 1
Color aware       : on, color-method (DEI)
acceptance        : CIR - FLR NA FTD 100 ms, FDV 10 ms
MTU               : 128, 256, 512,
CFM               : domain md000, ma ma10, mep 10
                    Destination e8c5.7adb.d59b
Stream s1, Test (cir-eir)
-----------------------------------------------------
Oper State        : Passed
Test Duration     : 00:02:03 (hh:mm:ss)
Start Time        : 2019 Feb 14 10:53:43
End Time          : 2019 Feb 14 10:55:46
CIR:
Tx packets 50720289, bytes 15249900226
Rx packets 25445881, bytes 8694009171
FL 25274408, FLR 0.0000
FD 11 us, FDV 1 us
FD Min 10 us, FD Max 34 us, FD exceeded 0%
Out of Order packets 0
Error packets 0
EIR:
Tx packets 25358536, bytes 7624466320
Rx packets 12722142, bytes 4346731850
FL 12636394, FLR 0.0000
FD 11 us, FDV 1 us
FD Min 10 us, FD Max 36 us, FD exceeded 0%
Out of Order packets 0
Error packets 0
OcNOS#
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
CBS	Committed Burst Size
CIR	Committed Information Rate
CM	Color Mode

CTF	Collector Traffic Function
DEI	Drop Eligible Indicator
DMM	Delay Measurement Message
DP	Drop Precedence
EBS	Excess Burst Size
EIR	Excess Information Rate
EMIX	Ethernet Mix
FDV	Frame Delay Variation
FLR	Frame Loss Ratio
FTD	Frame Transfer Delay
GTF	Generate Test Function
IR	Information Rate
KPI	Key Performance Indicator
LBM	Loopback Message
LBR	Loopback Reply
MAC	Medium Access Control
MP	Measurement Point
MTU	Maximum Transmission Unit
QoS	Quality of Service
SAC	Service Acceptance Criteria
SAT	Service Activation Test
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
ToD	Time of Day
ULR	Utilized Line Rate
UNI	User Network Interface
UNI-C	UNI – Customer
UNI-N	UNI – Network
VLAN	Virtual LAN