

ipinfusion™

OcNOS®

**Open Compute
Network Operating System
for Data Centers
Version 6.5.2**

Key Features

August 2024

© 2024 IP Infusion Inc. All Rights Reserved.

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

IP Infusion Inc.

3965 Freedom Circle, Suite 200

Santa Clara, CA 95054

+1 408-400-1900

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

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

support@ipinfusion.com

Trademarks:

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

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

Contents

Preface	5
Audience	5
Conventions	5
Related Documentation	5
Feature Availability	5
Migration Guide	5
Support	5
Comments	6
CHAPTER 1 PIM Sparse-Dense Mode Configuration	12
Overview	12
Feature Characteristics	12
Configuration	12
Topology	13
Enabling IP Multicast Routing	14
Enabling PIM-SMDM	15
Sparse Mode Operation versus Dense Mode Operation	16
CHAPTER 2 MLD Configuration	21
Overview	21
Feature Characteristics	21
MLD Versions	21
MLD Leave Operation	22
Configuration	23
Topology	23
MLD Configuration	23
Configuring MLD Parameters	24
MLD Group Table after MLDv1 Membership Report is received	26
Glossary	27
CHAPTER 3 MLD Snooping Configuration	28
Overview	28
Feature Characteristics	28
MLD Snooping Configuration	29
Glossary	32
CHAPTER 4 PIM Source-Specific Multicast Configuration	33
Overview	33
Feature Characteristics	33
PIM-SSM Configuration	33
Topology	34
Enable IP Multicast Routing on all Routers	34
CHAPTER 1 Streaming Telemetry Dial-Out Mode	42
Overview	42
Prerequisites	44
Configuration	44

Implementation Examples.....	49
Dial-Out Commands	49
Revised CLI Commands.....	59
Glossary.....	59
CHAPTER 2 DHCPv6 Prefix Delegation Configuration	61
Overview	61
Configuration.....	61
DHCP Multiple Prefix Delegation Command	67
Revised CLI Commands.....	67
Glossary.....	68
CHAPTER 3 EVPN VXLAN E-Tree	69
Overview	69
Prerequisites	70
Configuration.....	74
Implementation Examples.....	90
E-Tree CLI Commands	91
Revised CLI Commands.....	91
Troubleshooting	92
Glossary.....	92
CHAPTER 1 ISIS Multi Topology	95
Overview	95
Prerequisite.....	95
Configuration.....	95
Validation for Multi Topology	102
CLI Commands.....	136
Glossary.....	137

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
<code>monospaced type</code>	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.

Enhanced Security and Performance

This section describes the security, performance, scalability, and access control enhancements and new features introduced in the Release 6.5.2.

- [PIM Sparse-Dense Mode Configuration](#)
- [MLD Configuration](#)
- [MLD Snooping Configuration](#)
- [PIM Source-Specific Multicast Configuration](#)

CHAPTER 1 PIM Sparse-Dense Mode Configuration

Overview

Protocol Independent Multicast Sparse Mode-Dense Mode (PIM-SMDM) is a protocol designed to manage both sparse and dense multicast groups, efficiently handling varying multicast distribution patterns. In dense mode, it assumes listeners on all subnetworks, initially flooding the network and then pruning back areas without listeners. In sparse mode, it assumes few listeners and forwards traffic only to known listeners, reducing unnecessary transmission. PIM-SMDM switches between modes based on the multicast group's status, treating interfaces accordingly. A group is sparse if the router knows about a Rendezvous Point (RP) for it. Its adaptability makes it a versatile solution for diverse network scenarios.

Feature Characteristics

Protocol Independent Multicast Sparse Mode-Dense Mode (PIM-SMDM) manages both sparse and dense multicast groups, seamlessly switching modes based on the group's status. In dense mode, it floods the network with multicast traffic and prunes areas without listeners. In sparse mode, it forwards traffic only to known listeners, reducing unnecessary data transmission. The protocol treats interfaces as dense or sparse based on the group's mode, considering a group sparse if a Rendezvous Point (RP) is known. PIM-SMDM efficiently distributes multicast streams, optimizes network resources, and adapts to different multicast group modes, making it suitable for diverse network scenarios.

Benefits

- Manages both sparse and dense multicast groups simultaneously, making it adaptable to various network scenarios.
- Seamlessly switches between dense and sparse modes based on the multicast group's status, ensuring efficient distribution of multicast streams.
- Reduces unnecessary data transmission in sparse mode by forwarding traffic only to known listeners, optimizing the use of network resources.
- Can handle networks of different sizes and complexities, adapting to the number of listeners and the multicast group's distribution patterns.
- By pruning areas without listeners in dense mode and reducing traffic in sparse mode, PIM-SMDM enhances overall network performance and minimizes congestion.
- Treats interfaces as dense or sparse based on the group's mode, dynamically adapting to the network's current multicast requirements.
- Utilizes Rendezvous Points (RPs) in sparse mode for efficient centralized management of multicast sources and receivers.
- Suitable for a wide range of applications, from small-scale deployments to large, complex networks with varying multicast distribution needs.

Configuration

The required steps to configure PIM-SMDM are the following:

- Enable IP multicast on each PIM router (see [Enabling IP Multicast Routing](#))
- Enable PIM-SMDM on the desired interfaces (see [Enabling PIM-SMDM](#))
- Example for the group operating in sparse-mode having Static RP (see [Configuring Rendezvous Point Statically for PIM-SMDM](#))
- Example for the group operating in dense-mode having no RP

All multicast group states are dynamically maintained as the result of IGMP Report/Leave and PIM Join/Prune messages. This section provides the steps to configure the PIM-SMDM feature. Configuration steps and examples are used for two relevant scenarios. The following figure displays the network topology used in these examples.

Topology

There are two topologies for this sparse-dense mode configuration. Understanding these modes helps network administrators select the best multicast strategy for their network, ensuring efficient and reliable traffic delivery.

Sparse Mode

The network topology in [Figure 1-1](#), includes several routers and hosts within a multicast network. Key components are Router_C, the Rendezvous Point (RP), and Host_1 and Host_2, which join a multicast group. Subnet 1 connects Host_1, Host_2, Router_E, and Router_F, with the latter two managing multicast traffic on the subnet.

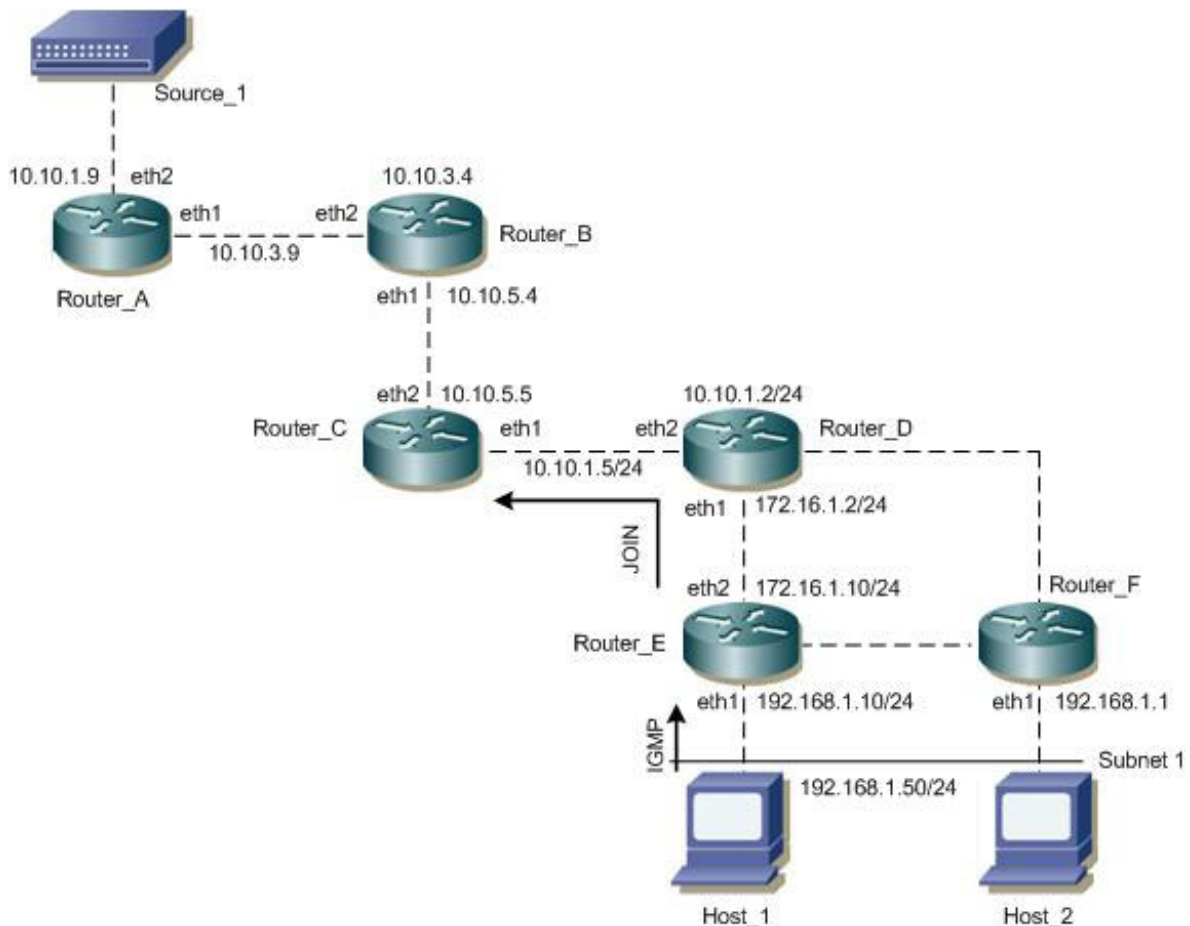


Figure 1-1: PIM-SMDM Configuration Topology (a)

Dense-mode

In the network topology shown in [Figure 1-2](#), Source_1 (10.10.1.52) sends multicast data to group address 224.0.1.3. Host_1 shows interest in this group by sending an IGMP membership report, which Router_C processes to associate its eth1 interface with the group. As data packets flow from Source_1, each router creates an (S,G) entry in its multicast routing table. Router_C forwards the packets through eth1 to Host_1 and, having a downstream receiver, does not send a prune message to its upstream neighbor, Router_E, ensuring continuous delivery of multicast traffic to interested hosts.




Figure 1-2: PIM-SMDM Configuration Topology (b)

Enabling IP Multicast Routing

To enable IP multicast routing on all of the PIM routers inside the PIM domain:

1. Enter the config mode.
R1(config)#configure terminal
2. Enable the IP Multicast routing.
(config)#ip multicast-routing
3. To commit the changes and exit.
(config)#commit
(config)#exit

Enabling PIM-SMDM

Enable PIM-SMDM on all participating interfaces within each of routers inside the PIM domain on which you want to run PIM. In the following sample configuration, both eth1 and eth2 are enabled for PIM-SMDM on the router.

1. Enter the config mode.
R1(config)#configure terminal
2. Configure the interface (eth1).
(config)#interface eth1
(config)#ip pim sparse-dense-mode
3. To commit the changes and exit.
(config)#commit
(config)#exit
4. Configure interface (eth2).
(config)#interface eth2
(config)#ip pim sparse-dense-mode
5. To commit the changes and exit.
(config)#commit
(config)#exit

Validation

Here is the sample configuration for Router_C:

```
hostname Router_C
!
interface eth0
!
interface eth1
 ip pim sparse-dense-mode
!
interface eth2
 ip pim sparse-dense-mode
!
interface lo
!
!
ip multicast-routing
```

The show ip pim interface command displays the interface details for Router_C.

```
Router_C#show ip pim interface
```

Address	Interface	VIFindex	Ver/ Mode	Nbr Count
192.168.1.10	eth1	0	v2/SD	0
172.16.1.10	eth2	2	v2/SD	1

Sparse Mode Operation versus Dense Mode Operation

The following examples differentiates the group operating in sparse mode versus dense mode:

- Sparse mode operation when the RP is present for the group
- Dense mode operation when there is no RP for the group.

Sparse Mode Operation

Configuring Rendezvous Point Statically for PIM-SMDM

Every PIM multicast group needs to be associated with the IP address of a Rendezvous Point (RP), which is a router that resides in a multicast network domain. The address of the RP is used as the root of a group-specific distribution tree. All nodes in the domain that want to receive traffic sent to the group are aware of the address of the RP. For all senders to reach all receivers within a group, all routers in the domain must be able to map to the RP address configured for the group. There can be several RPs configured in a network deploying PIM-SM, each serving a different group.

You can statically configure a RP by specifying the RP address in every router in the PIM domain. The use of statically configured RPs is ideal for small network environments or ones that do not require many RPs and/or require changing the assignment of the RPs often. Changing the assignment of an RP requires the re-configuration of the RP address in all of the routers in the PIM domain.

In static RP configurations, RP failover is not available.

When configuring the RP statically, do the following:

- On every router, include the `ip pim rp-address A.B.C.D` statement even if a router does not have any source or group member attached to it
- Assign only one RP address for a multicast group in the PIM domain

The network topology shown in the [Figure 1-1](#), includes several routers, a source, and hosts in different subnets.

- Source_1:
 - Connected to Router_A via eth2 with IP address 10.10.1.9.
- Router_A:
 - Interface eth1 connects to eth2 of Router_B with IP address 10.10.3.9.
- Interface eth2 connects to Source_1.
- Router_B:
 - Interface eth1 connects to eth2 of Router_A with IP address 10.10.3.4.
 - Interface eth2 connects to eth1 of Router_C with IP address 10.10.5.4.
- Router_C:
 - Interface eth1 connects to eth2 of Router_B with IP address 10.10.5.4.
 - Interface eth2 connects to eth1 of Router_D with IP address 10.10.5.5 and 10.10.1.5/24 network.
- Router_D:
 - Interface eth1 connects to eth2 of Router_C with IP address 10.10.1.2/24.
 - Interface eth2 connects to eth1 of Router_E with IP address 172.16.1.2/24.
- Router_E:
 - Interface eth1 connects to eth2 of Router_D with IP address 172.16.1.2/24.

- Interface eth2 connects to eth1 of Router_F with IP address 172.16.1.10/24.
- Interface eth1 connects to Host_1 via IGMP with IP address 192.168.1.10/24.
- Router_F:
 - Interface eth1 connects to eth2 of Router_E with IP address 172.16.1.10/24.
 - Interface eth1 connects to Host_2 with IP address 192.168.1.50/24.
- Host_1:
 - Connected to Router_E with IP address 192.168.1.10/24.
- Host_2:
 - Connected to Router_F with IP address 192.168.1.50/24.

Configure Static RP

1. Enter the config mode.
`R1(config)#configure terminal`
2. Configure interface (eth1).
`(config)#ip pim rp-address 10.10.1.5`
3. To commit the changes and exit.
`(config)#commit`
`(config)#exit`

Validation

Here is the sample configuration for Router_D:

```
hostname Router_D
!
interface eth0
!
interface eth1
 ip pim sparse-dense-mode
!
interface eth2
 ip pim sparse-dense-mode
!
interface lo
!
!
ip multicast-routing
ip pim rp-address 10.10.1.5
!
```

RP Details

At Router_D, the `show ip pim rp mapping` command shows that 10.10.1.5 is the RP for all multicast groups 224.0.0.0/4, and is statically configured. All other routers will have a similar output:

```
Router_D#sh ip pim rp mapping
PIM Group-to-RP Mappings
Override RP cnt: 0
```

```
Group(s): 224.0.0.0/4, Static
  RP: 10.10.1.5
    Uptime: 00:01:45
```

At Router_D, use the `show ip pim rp-hash` command to display the selected RP for a specified group (224.0.1.3):

```
Router_D#show ip pim rp-hash 224.0.1.3
RP: 10.10.5.37
```

Interface Details

The `show ip pim interface` command displays the interface details for Router_E, and shows that Router_E is the Designated Router on Subnet 1.

```
Router_E#show ip pim interface
Address          Interface VIFindex Ver/   Nbr    DR    DR
                  Mode     Count   Prior
192.168.1.10     eth1     0       v2/SD  1      1     192.168.1.10
172.16.1.10      eth2     2       v2/SD  1      1     172.16.1.10
```

IP Multicast Routing Table

Note: The multicast routing table displays for an RP router are different from other routers.

The `show ip pim mroute` command displays the IP multicast routing table. In this table, the following fields are defined:

- RPF nbr Displays the unicast next-hop to reach RP.
 and mask length.
- RPF idx Displays the incoming interface for this (*, G) state.
- RP Displays the IP address for the RP router
- B Displays the bidirectional pim mode

The leading dots
 Stand for VIF index

```
Router_E#show ip pim mroute
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 0
(S,G,rpt) Entries: 0
(*, 224.0.1.3)
RP: 10.10.1.5
RPF nbr: 172.16.1.2
RPF idx: eth2
Upstream State: JOINED
Local      .....
Joined    j.....
Asserted  .....
Outgoing  o.....
```

At Router_E, eth2 is the incoming interface of the (*, G) entry, and eth1 is on the outgoing interface list of the (*, G) entry. This means that there is a group member through eth1, and the RP is reachable through eth2.

The 0 position on this 32-bit index is for eth1 (as illustrated in the interface display above). The j on the 0 index indicates that the Join has come from eth1.

Since Router_C is the RP, and the root of this multicast tree, the `show ip pim mroute` command on Router_C shows RPF nbr as 0.0.0.0 and RPF idx as none.

```
Router_C#show ip pim mroute
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 0
(S,G,rpt) Entries: 0
(*, 224.0.1.3)
RP: 10.10.1.5
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: JOINED
Local      .....
Joined     j.....
Asserted   .....
Outgoing   o.....
```

For configuring Rendezvous point dynamically refer [Configure Rendezvous Point Dynamically Using Bootstrap Router Method](#) and [Configuring Rendezvous Point Statically](#)

Dense-mode Operation

The network topology described in [Figure 1-2](#), the Source_1 address is 10.10.1.52 and the group address is set to 224.0.1.3.

In this example all routers are running PIM-SMDM.

- Host_1 sends an IGMP membership report to Subnet 1.
- After Router_C receives this report, it associates its receiving interface, eth1, with the group reported in the IGMP message, for example, group1.
- Source_1 then sends a data packet for group1.
- Every router creates an (S,G) entry in the multicast routing table.
- When the data packet reaches Router_C, it forwards via the interface, eth1, because there is a local member on this interface for this group. Router_C has a downstream receiver, so it does not send a prune message to its upstream neighbor router, Router_E.

The network topology shown in the [Figure 1-2](#), includes a source, three routers, and a host in a subnet.

- Source_1:
 - Connected to Router_A via eth2 with IP address 10.10.1.52 and sending multicast traffic to the multicast group 224.0.1.3.
- Router_A:
 - Interface eth1 connects to eth2 of Router_B with IP address 10.10.3.9.
 - Interface eth2 connects to Source_1 with IP address 10.10.1.9.
- Router_B:
 - Interface eth1 connects to eth2 of Router_A with IP address 10.10.3.4.

- Interface eth2 connects to eth1 of Router_C with IP address 172.16.1.2/24.
- Router_C:
 - Interface eth1 connects to eth2 of Router_B with IP address 172.16.1.2/24.
 - Interface eth2 connects to Host_1 via eth1 with IP address 172.16.1.10/24.
 - Interface eth1 connects to Host_1 with IP address 192.168.1.10/24 and 192.168.1.50/24 via IGMP.
- Host_1:
 - Connected to Router_C with IP address 192.168.1.50/24 and subscribed to the multicast group 224.0.1.3 via IGMP.

Validation

Enter the commands listed in this section to confirm the previous configurations.

IP Multicast Routing Table

The `show ip pim mroute` command displays the IP multicast routing table.

```
Router_C#show ip mroute
IP Multicast Routing Table
Flags: I - Immediate Stat, T - Timed Stat, F - Forwarder installed
Timers: Uptime/Stat Expiry Interface State:
Interface (TTL) (10.10.1.52, 224.0.1.3), uptime 00:00:15
Owner PIM-DM, Flags: F
Incoming interface: eth2
Outgoing interface list:
eth1 (1)
```

IP PIM-SMDM Multicast Routing Table

The `show ip pim dense-mode mroute` command displays the IP PIM-DM multicast routing table

```
Router_C#show ip pim mroute
PIM-DM Multicast Routing Table (10.10.1.52, 224.0.1.3)
RPF Neighbor: 172.16.1.2, Nexthop: 172.16.1.2, eth2
Upstream IF: eth2
Upstream State: Forwarding
Assert State: NoInfo
Downstream IF List:
eth1, in 'olist': Downstream State: NoInfo Assert State: NoInfo
```

CHAPTER 2 MLD Configuration

Overview

Multicast Listener Discovery (MLD) is a protocol used by IPv6 hosts to communicate their desire to receive multicast traffic to the neighboring multicast routers. It serves a role similar to that of Internet Group Management Protocol (IGMP) in IPv4 networks. MLD is essential for efficient multicast routing in IPv6 networks, ensuring that multicast data is only sent to network segments with interested receivers.

IP hosts use MLD to inform multicast routers about their membership in specific multicast groups, allowing routers to maintain a list of group memberships per interface. When a host joins a multicast group, it sends an MLD Report to the router, which updates its membership list. Routers then use this information to forward multicast data only to network segments with interested hosts, optimizing network resources by preventing unnecessary traffic.

By default, when PIMv6 is enabled on an interface, MLD version 2 is enabled. MLD can be enabled on an interface explicitly.

Feature Characteristics

MLD allows hosts to notify multicast routers about their interest in joining or leaving multicast groups, with routers maintaining membership lists for each interface. Hosts use MLD Report messages to join and Done messages to leave groups, enabling routers to update memberships. Routers then use this data to forward multicast traffic only to interested network segments, optimizing bandwidth. By default, MLDv2 is enabled with PIMv6, supporting source-specific multicast and maintaining compatibility with MLDv1. Administrators can manually configure MLD on interfaces as needed, ensuring effective multicast management and interoperability between versions.

Benefits

These benefits make MLD an essential protocol for efficient and effective multicast routing in IPv6 networks, enhancing performance, scalability, and resource utilization.

- Efficient Multicast Traffic Management
- Network Resource Optimization
- Improved Scalability
- Enhanced Performance and Reliability
- Compatibility and Interoperability
- Administrative Control and Flexibility.

MLD Versions

OcNOS supports MLDv1 and MLDv2. By default, OcNOS enables MLDv2 when PIMv6 is enabled on an interface.

MLDv2 includes the following key changes from MLDv1:

- Support for Source-Specific Multicast (SSM), which builds shortest path trees from each receiver to the source, through the following feature:
 - Host messages that can specify both the group and the source.

- The multicast state that is maintained for groups and sources, not just for groups as in MLDv1.
- Hosts no longer perform report suppression, which means that hosts always send MLD membership reports when an MLD query message is received.

MLD Operation

MLD works on the premise of three major packets exchange between MLD enabled routers and hosts, interested in joining a particular group.

MLD Query Operation

Once MLD is enabled or PIMv6 is enabled (which enables MLDv2), on any interface it starts sending Query message, which is called general query to the all-hosts multicast group at ff02::1 periodically to discover whether any hosts want to receive multicast data.

OcNOS elects a router as the MLD querier on a subnet if it has the lowest IP address. As long as a router continues to receive query messages from a router with a lower IP address, it resets a timer that is based on its querier timeout value. If the querier timer of a router expires, it becomes the designated querier. If that router later receives a host query message from a router with a lower IP address, it drops its role as the designated querier and sets its querier timer again.

In the figure below Router-1 eth2 sends query every query-interval. Since Router1-eth2 IPv6 link local address is less than Router-2 eth2, Router-1 eth2 becomes querier on the LAN.

MLD Membership Report Operation

When a host receives a query from the local router it sends a Host Membership Report for all the multicast groups for which it wants to receive multicast traffic. This is called solicited membership report.

When a host joins a new group, the host immediately sends a Membership Report to inform a local router that it wants to receive multicast traffic for the group it has just joined without waiting to receive a Query. This is called unsolicited membership report.

In the figure below Host-1 and Host-2 sends membership reports to Router-1 eth2 for all the multicast groups for which they want to receive multicast traffic. Upon reception of membership report Router-1 maintains an MLD group table containing multicast group-address, interface name on which it receives the report.

MLD Leave Operation

When a multicast host leaves a group, a host that runs MLD sends an MLD leave message. To check if this host is the last host to leave the group, the router sends an MLD query (Called as Group-specific-query) message and starts a timer that you can configure called the last member query response interval. If no reports are received before the timer expires, the software removes the group state. The router continues to send multicast traffic for a group until its state is removed.

In the figure below Host-1 and Host-2 sends leave message to Router-1 eth2 for all the multicast groups for which they don't want to receive multicast traffic. In response to leave message Router-1 eth2 sends an group-specific-query message before removing the multicast group address from the MLD table.


Configuration

You can configure MLD on a network device to manage multicast group memberships effectively. This configuration enables efficient multicast traffic distribution, optimizes bandwidth usage, and ensures that multicast data is only sent to network segments with interested receivers.

Topology

This topology ensures that each router's interfaces are configured with the specified IP or IPv6 link-local addresses, and verifies the switch's configurations for connectivity. It involves setting up routing protocols or static routes on each router for communication, and assigning and configuring IPv6 addresses on router and host interfaces to ensure proper device communication via link-local addresses. Additionally, routers are configured to handle unicast or multicast traffic, with necessary multicast routing protocols set up for multicast traffic.

The network topology shown in the [Figure 2-3](#) includes three routers, a switch, two hosts, and a source.



MLD Configuration

To configure Multicast Listener Discovery (MLD) on a network device, follow these steps:

1. Enter the config mode.

```
R1(config)#configure terminal
```

2. Enable IPv6 Multicast routing.

```
(config)#ipv6 multicast-routing
```

3. Configure the interface.

```
(config)#interface eth2
```

4. Assign the IP address to the interface.

```
(config-if)#ip address 2001::1/64
```

5. Enable MLD version 1 on the interface.

```
(config-if)#ipv6 mld version 1
```

6. To commit the changes and exit.

```
(config)#commit
```

```
(config)#exit
```

Validation

Enter the commands listed in this section to confirm the previous configurations.

```
#show running-config
!
no service password-encryption
!
hostname rtr1
!
Ipv6 multicast-routing
!
!
interface eth2
ip address 2001::1/64
no shutdown
ipv6 mld version 1
```

Configuring MLD Parameters

The configuration that follows shows how MLD parameters can be configured.

1. Enter the config mode.

```
R1(config)#configure terminal
```

2. Configure the interface for MLD, enter interface configuration mode for eth2

```
(config)#interface eth2
```

3. Assign the IPv6 address to the interface.

```
(config-if)#ip address 2001::1/64
```

4. Enable MLD version 1.

```
(config-if)#ipv6 mld version 1
```

5. Enable IPv6 Multicast routing, execute the following command to enable IPv6 multicast routing.

```
(config)#ipv6 multicast-routing
```

6. Apply MLD Configuration to the Interface, configure the MLD access group.

```
(config)#interface eth2
(config-if)#ipv6 mld access-group 1
```

7. Enable MLD immediate leave.

```
(config-if)#ipv6 mld immediate-leave
```

8. Configure MLD group-list.

```
group-list 1
```

9. Set the last member query count.

```
(config-if)# ipv6 mld last-member-query-count 7
```

10. Set the last member query interval.

```
(config-if)# ipv6 mld last-member-query-interval 25500
```

11. Limit the number of MLD groups.

```
(config-if)#ipv6 mld limit 100
```

12. Set the MLD querier timeout, interval, query maximum response time, robustness variable, startup query count, and startup query interval.

```
(config-if)#ipv6 mld querier-timeout 300
(config-if)#ipv6 mld query-interval 200
(config-if)#ipv6 mld query-max-response-time 150
(config-if)#ipv6 mld robustness-variable 4
(config-if)#ipv6 mld startup-query-count 4
(config-if)# ipv6 mld startup-query-interval 50
(config-if)#ipv6 mld static-group FF1E::1
```

13. To commit the changes and exit.

```
(config)#commit
(config)#exit
```

Validation

Enter the commands listed in this section to confirm the previous configurations

```
#show running-config
!
no service password-encryption
!
hostname rtl1
!
!
Ipv6 multicast-routing
!
!
interface eth2
ipv6 address 2001::1/64
no shutdown
ipv6 mld access-group 1
```

```

ipv6 mld immediate-leave group-list 1
ipv6 mld last-member-query-count 7
ipv6 mld limit 100
ipv6 mld static-group ffile::1
ipv6 mld last-member-query-interval 25500
ipv6 mld querier-timeout 300
ipv6 mld query-interval 200
ipv6 mld query-max-response-time 150
ipv6 mld startup-query-interval 50
ipv6 mld startup-query-count 4
ipv6 mld robustness-variable 4
ipv6 mld ra-option
ipv6 mld version 1
!!

```

```

Rtr1#show ipv6 mld interface eth2
Interface eth2 (Index 4)
MLD Enabled, Active, Querier, Configured for version 1
Internet address is fe80::eac5:7aff:fe78:a2d9
MLD interface limit is 100
MLD interface has 1 group-record states
MLD interface statistics:
v1-reports: 0, v1-leaves: 0
v2-reports: 0
MLD query interval is 200 seconds
MLD startup query interval is 50 seconds
MLD startup query count is 4
MLD querier timeout is 300 seconds
MLD max query response time is 150 seconds
Group Membership interval is 950 seconds
MLD Last member query count is 7
Last member query response interval is 1000 milliseconds

```

MLD Group Table after MLDv1 Membership Report is received

MLD group table is populated at router by virtue of either static join is configured on interface or dynamic report is being received on the interface.

The `show ipv6 mld groups` command displays the MLD group table. In this table, the following fields are defined.

Table 2-1: MLD group table after MLDv1 membership report

Group address	Displays the Multicast Group for which report is received.
Interface	Interface name on which Membership report is received.
Uptime	Duration since the report is received.
Expiry	Time frame in which the multicast group is going to expire.
Last Reporter	Host address from where the report is generated.

```

#show ipv6 mld groups
MLD Connected Group Membership

```

Group Address	Interface	Uptime	Expires	State	Last Reporter
ff04::1	xe18	00:00:10	00:15:40	Active	fe80::1
ffle::1	xe18	00:17:22	static	Active	::

```
#show ipv6 mld groups detail
MLD Connected Group Membership Details
```

```
Flags: (M - SSM Mapping, R - Remote,
        SG - Static Group, SS - Static Source)
```

```
Interface:      xe18
Group:          ff04::1
Flags:          R
Uptime:         00:00:33
Group mode:     Exclude (Expires: 00:15:17)
State:          Active
Last reporter:  fe80::1
Source list is empty
```

```
Flags: (M - SSM Mapping, R - Remote,
        SG - Static Group, SS - Static Source)
```

```
Interface:      xe18
Group:          ffle::1
Flags:          SG
Uptime:         00:17:45
Group mode:     Exclude (Static)
State:          Active
Last reporter:  ::
Source list is empty
```

Glossary

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

Key Terms/Acronym	Description
IGMP	Multicast Listener Discovery (MLD) is a protocol used in IPv6 networks that allows network devices (hosts) to inform multicast routers of their intention to receive multicast traffic.
MLD	The Internet Group Management Protocol (IGMP) is a communication protocol used in IPv4 networks to manage multicast group memberships.

CHAPTER 3 MLD Snooping Configuration

Overview

In IPv6 networks, Multicast Listener Discovery (MLD) Snooping plays a crucial role in optimizing multicast traffic management within Layer-2 switches. By default, without MLD, Layer-2 switches treat IPv6 multicast traffic like broadcast traffic, forwarding frames received on one interface to all others. This indiscriminate forwarding leads to unnecessary traffic across the network, impacting performance.

MLD Snooping addresses this issue by intelligently monitoring and managing multicast traffic. Here's how it works: switches enabled with MLD Snooping analyze MLD messages exchanged between IPv6 hosts and multicast routers. Instead of flooding multicast traffic to all ports, switches learn which ports have hosts interested in specific multicast groups. They then selectively forward multicast traffic only to those ports where the interested hosts reside, significantly reducing network congestion and improving efficiency.

To enable MLD Snooping, administrators typically use the `switchport` command on each switch port to switch it to Layer-2 mode, allowing the switch to monitor MLD messages effectively. This approach ensures that multicast traffic is delivered only to the intended recipients, optimizing network performance and resource utilization in IPv6 environments.

Feature Characteristics

MLD Snooping enables Layer-2 switches to intelligently manage IPv6 multicast traffic by forwarding packets only to ports with active listeners for specific multicast groups, preventing unnecessary network-wide flooding. By selectively forwarding multicast traffic based on MLD messages exchanged between hosts and routers, MLD Snooping enhances overall network performance, reducing congestion and optimizing bandwidth usage. It eliminates broadcast-like behavior by maintaining a multicast group table and forwarding traffic solely to ports where interested hosts are located, akin to IPv4's IGMP Snooping. This efficient management conserves network resources, delivering packets only where there are active receivers, and reduces control plane overhead by handling just one MLD membership report per multicast group, even with multiple interested hosts.

Benefits

- Efficient Multicast Traffic Management
- Improved Network Performance
- Reduced Broadcast-Like Behavior
- Optimized Resource Utilization
- Reduced Control Plane Overhead
- Enhanced Security Features
- Compatibility and Integration.

Topology

In this topology, switch S1 configures eth1 as a multicast router port. Since MLD Snooping manages multicast traffic in bridged LAN setups, router R1 does not need to run MLD Snooping and can instead utilize any multicast protocol like PIMv6-SM. Therefore, this example focuses solely on configuring switch S1, and does not cover configuration details for router R1.




Figure 3-4: MLD Snooping Topology

As a result of this configuration:

- The switch itself replies with membership report messages in response to queries received on interface eth1. However, if you do not enable report suppression on the switch, when it receives an MLD Query message on eth1, it forwards it to both Host A and Host B. As a result, both hosts reply with a Membership report (as Layer-2 MLD is running on the hosts).
- Because Host A and Host B are members of the same multicast group, the router is not notified when A leaves the group, because the group still has another member. When Host B leaves the group, the switch will send a Leave message to the Router with the destination address as FF02::2(All Router Destination Address).

MLD Snooping Configuration

To enable MLD Snooping on an interface:

1. Add a bridge to the spanning-tree table
2. Specify the interface to be configured
3. Associate the interface with bridge group
4. MLD snooping will be enabled by default

5. Configure ports that are connected to routers as multicast router ports

6. By default, MLD report suppression is enabled on the switch

Note: Execute `I2 unknown mcast` CLI to enable the option to drop the unknown multicast traffic.

S1

1. Enter the config mode.

```
R1(config)#configure terminal
```

2. Enable the MLD on interface, set the bridge protocol and configure interface eth0.

```
(config)#bridge 1 protocol ieee vlan-bridge
```

```
(config)#interface eth0
```

```
(config-if)#shutdown
```

```
(config-if)#switchport
```

```
(config-if)#bridge-group 1
```

```
(config-if)#switchport mode access
```

```
(config-if)#no shutdown
```

```
(config-if)#exit
```

3. Configure interface eth1.

```
(config)#interface eth1
```

```
(config-if)#shutdown
```

```
(config-if)#switchport
```

```
(config-if)#bridge-group 1
```

```
(config-if)#switchport mode access
```

```
(config-if)#no shutdown
```

```
(config-if)#exit
```

4. Configure interface eth2.

```
(config)#interface eth2
```

```
(config-if)#shutdown
```

```
(config-if)#switchport
```

```
(config-if)#bridge-group 1
```

```
(config-if)#switchport mode access
```

```
(config-if)#no shutdown
```

```
(config-if)#exit
```

5. Configure interface vlan1.1 for MLD snooping.

```
(config)#interface vlan1.1
```

```
(config-if)# MLD snooping mrouter interface eth1
```

6. To commit the changes and exit.

```
(config)#commit
```

```
(config)#exit
```

Validation

```
#show running-config interface eth0
!
interface eth0
switchport
bridge-group 1
switchport mode access
!
#show running-config interface eth1
!
interface eth1
switchport
bridge-group 1
switchport mode access
!

#show running-config interface eth2
!
interface eth2
switchport
bridge-group 1
switchport mode access
!

#show mld snooping groups
MLD Snooping Group Membership
Group source list: (R - Remote, S - Static, > - Hw Installed)
Vlan  Group/source Address          Interface  Flags  Uptime
Expires  Last Reporter          Version
1       ff06::2                eth0      R      >    00:00:41
00:03:39 fe80::1                V2

#show mld snooping interface vlan1.1

MLD Snooping information for vlan1.1 (Index 25001)
MLD Snooping is globally enabled
MLD Snooping is enabled on this interface
MLD Active, Non-Querier,
MLD querying router is :
      :fe80::eac5:7aff:feb1:6b2d
MLD query interval is 125 seconds
MLD querier timeout is 255 seconds
MLD max query response time is 10 seconds
Last member query response interval is 1000 milliseconds
Group Membership interval is 260 seconds
MLD Snooping fast-leave is not enabled
MLD Snooping querier is not enabled
MLD Snooping report suppression is enabled
Number of Groups: 1
Number of v1-reports: 0
Number of v1-leaves: 0
Number of v2-reports: 3
Active Ports:
```

```
eth0
eth1
eth2
```

Glossary

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

Key Terms/Acronym	Description
MLD	The Internet Group Management Protocol (IGMP) is a communication protocol used in IPv4 networks to manage multicast group memberships.

CHAPTER 4 PIM Source-Specific Multicast Configuration

Overview

PIM Source-Specific Multicast (SSM) is a multicast routing protocol that enhances the efficiency and security of multicast communication by enabling hosts to receive multicast traffic directly from specific sources. Here's a detailed overview of how PIM SSM operates using a subset of PIM sparse mode and IGMPv3/MLDv2:

SSM utilizes PIM sparse mode (PIM-SM) to create a Shortest Path Tree (SPT) directly between multicast sources and receivers. Hosts signal their interest using IGMPv3 (IPv4) or MLDv2 (IPv6), specifying the source IP address to join multicast groups without requiring a Rendezvous Point (RP). This direct communication approach optimizes multicast efficiency by bypassing the RP and establishing efficient data paths tailored to specific source-receiver relationships, enhancing network performance and security in multicast environments.

PIM Source-Specific Multicast (SSM) thus enhances multicast communication by streamlining the process of delivering multicast traffic directly from sources to receivers, leveraging existing multicast protocols and minimizing network complexity.

Feature Characteristics

PIM SSM enables hosts to specify source IP addresses when joining multicast groups, facilitating direct communication paths and eliminating the need for a Rendezvous Point (RP). It leverages PIM sparse mode to establish efficient Shortest Path Trees (SPTs) between sources and receivers, ensuring optimized multicast traffic delivery. Hosts use IGMPv3 (IPv4) and MLDv2 (IPv6) for precise membership management, enhancing network security and efficiency by reducing unnecessary traffic and simplifying configuration. SSM supports scalable deployment alongside existing multicast infrastructure, promoting interoperability and streamlined network administration while optimizing resource utilization and improving overall network reliability.

Benefits

The benefits of PIM SSM:

- Efficient Multicast Traffic Handling
- Optimized Resource Utilization
- Enhanced Security
- Simplified Configuration and Management
- Scalability and Compatibility
- Improved Network Performance
- Support for Diverse Applications.

PIM-SSM Configuration

The required steps to configure PIM-SSM are the following:

- Enable IP multicast on each PIM router (see Enabling IP Multicast Routing)
- Enable PIM-SM on the desired interfaces (see Enable PIM-SM on an Interface)

- Configure PIM-SSM on router.

All multicast group states are dynamically maintained as the result of IGMP Report/Leave and PIM Join/Prune messages.

Topology

The following figure displays the network topology used in these examples.




Figure: PIMV6-SSM Configuration Topology
Figure 4-5: PIM-SSM Configuration Topology

Enable IP Multicast Routing on all Routers

1. Enter the config mode.
R1(config)#configure terminal
2. Enable the IP Multicast routing.
(config)#ip multicast-routing
3. Enable the IPv6 Multicast routing.
(config)#ipv6 multicast-routing

4. To commit the changes and exit.

```
(config)#commit
```

```
(config)#exit
```

Enable PIM SSM Default on all Routers

1. Enter the config mode.

```
R1(config)#configure terminal
```

2. Enable the PIM SSM for IPv4.

```
(config)# ip pim ssm default
```

3. Enable the PIM SSM for IPv6.

```
(config)# ipv6 pim ssm default
```

4. To commit the changes exit.

```
(config)#commit
```

```
(config)#exit
```

Enable PIM-SSM configuration on Router A

In the following sample configuration, both eth1 and eth2 are enabled for PIM-SSM on the router.

1. Enter the config mode.

```
R1(config)#configure terminal
```

2. Enable PIM-SSM configuration on router A, configure Interface eth1 and eth2.

```
(config)#interface eth1
```

```
(config-if)#ip address 10.1.1.1/24
```

```
(config-if)#ipv6 address 001::1/64
```

```
(config-if)#ip pim sparse-mode
```

```
(config-if)#ipv6 pim sparse-mode
```

```
(config-if)#ip igmp version 3
```

```
(config-if)#ipv6 mld version 2
```

```
(config-if)#exit
```

```
(config)#interface eth2
```

```
(config-if)#ip address 100.1.1.1/24
```

```
(config-if)#ipv6 address 2001::1/24
```

```
(config-if)#ip pim sparse-mode
```

```
(config-if)#ipv6 pim sparse-mode
```

```
(config-if)#ip igmp version 3
```

```
(config-if)#ipv6 mld version 2
```

3. To commit the changes and exit.

```
(config)#commit
(config)#exit
```

Enable PIM-SSM configuration on Router B

In the following sample configuration, both eth1 and eth2 are enabled for PIM-SSM on the router.

1. To enter the config mode, execute the following command.

```
R1(config)#configure terminal
```

2. Enable PIM-SSM configuration on router B, configure Interface eth2 and eth1.

```
(config)#interface eth2
(config-if)#ip address 10.1.1.2/24
(config-if)#ipv6 address 3001::2/64
(config-if)#ip pim sparse-mode
(config-if)#ipv6 pim sparse-mode
(config-if)#ip igmp version 3
(config-if)#ipv6 mld version 2
(config-if)#exit
(config)#interface eth1
(config-if)#ip address 11.1.1.1/24
(config-if)#ipv6 address 4001::1/24
(config-if)#ip pim sparse-mode
(config-if)#ipv6 pim sparse-mode
(config-if)#ip igmp version 3
(config-if)#ipv6 mld version 2
```

3. To commit the changes and commit.

```
(config)#commit
(config)#exit
```

Enable PIM-SSM configuration on Router C

In the following sample configuration, both eth1 and eth2 are enabled for PIM-SSM on the router.

1. To enter the config mode, execute the following command

```
R1(config)#configure terminal
```

2. Enable PIM-SSM configuration on router C, configure Interface eth2 and eth1.

```
(config)#interface eth2
(config-if)#ip address 11.1.1.2/24
(config-if)#ipv6 address 4001::2/64
(config-if)#ip pim sparse-mode
```



```
(config-if)#ipv6 pim sparse-mode
(config-if)#ip igmp version 3
(config-if)#ipv6 mld version 2
(config-if)#exit
(config)#interface eth1
(config-if)#ip address 101.1.1.1/24
(config-if)#ipv6 address 5001::1/24
(config-if)#ip pim sparse-mode
(config-if)#ipv6 pim sparse-mode
(config-if)#ip igmp version 3
(config-if)#ipv6 mld version 2
```

3. To commit the changes and commit.

```
(config)#commit
(config)#exit
```

Validation

Enter the commands listed in this section to confirm the previous configurations.

Interface Details

The show ip pim interface command displays the interface details for Router_C, and shows that Router_C is the Designated Router on Subnet 1.

```
Router_C#show ip pim interface
Address          Interface VIFindex Ver/   Nbr    DR    DR
                  Mode     Count  Prior
192.168.1.10     eth1     0      v2/S   1      1     192.168.1.10
172.16.1.10      eth2     2      v2/S   1      1     172.16.1.10
```

```
ROUTER C#show ipv6 pim interface
Total number of PIM interfaces:2
Interface VIFindex Ver/   Nbr    DR
                  Mode     Count  Prior
eth2          0      v2/D   1      1
  Address      : fe80::eac5:7aff:fea8:7cb9
  Global Address: 3001::1
eth1          1      v2/D   0      1
  Address      : fe80::eac5:7aff:fea8:7cc3
  Global Address: 2001::1
```

```
ROUTER C#sh ipv6 pim neighbor
```

Total number of PIM neighbors:2

```
Neighbor Address          Interface  Uptime/Expires  DR
```

			Pri/Mode
fe80::eac5:7aff:fea8:7cb9	eth1	01:29:52/00:01:18	1 /
fe80::eac5:7aff:feb1:6b13	eth2	01:29:49/00:01:28	1 /

Validation on IP Multicast Routing Table

Note: The multicast routing table displays for an S,G entries.

The show ip pim mroute command displays the IP multicast routing table. In this table, the following fields are defined:

```
LHR#sh ip pim mroute
IP Multicast Routing Table

(*,*,RP) Entries: 0
G/prefix Entries: 0
(*,G) Entries: 0
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

(101.1.1.2, 239.1.1.1)
RPF nbr: 10.1.1.2
RPF idx: xe14
SPT bit: 1
Upstream State: JOINED
Local      ..i.....
Joined     .....
Asserted   .....
Outgoing   ..o.....

(101.1.1.2, 239.1.1.1, rpt)
RP: 0.0.0.0
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: RPT NOT JOINED
Local      .....
Pruned     .....
Outgoing   .....
```

```
LHR#sh ipv6 pim mroute
IPv6 Multicast Routing Table

(*,*,RP) Entries: 0
G/prefix Entries: 0
(*,G) Entries: 0
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

(5001::2, ff06::2)
RPF nbr: fe80::36ef:b6ff:fe94:3ddd
```

```
RPF idx: xe14
SPT bit: 0
Upstream State: JOINED
  Local      ..i.....
  Joined     .....
  Asserted   .....
  Outgoing   ..o.....
```

(5001::2, ff06::2, rpt)

```
RP: ::
RPF nbr: ::
RPF idx: None
Upstream State: RPT NOT JOINED
  Local      .....
  Pruned     .....
  Outgoing   ..o.....
```

The ip igmp group detail and ipv6 mld group detail shows the source included (SSM)

LHR#sh ip igmp groups

```
IGMP Instance wide G-Recs Count is: 1
IGMP Connected Group Membership
```

Group Address	Interface	Uptime	Expires	State	Last Reporter
239.1.1.1	xe26	00:00:26	stopped	Active	100.1.1.2

LHR#sh ip igmp groups detail

```
IGMP Instance wide G-Recs Count is: 1
IGMP Connected Group Membership Details
```

Flags: (M - SSM Mapping, R - Remote, L - Local, SG - Static Group, SS - Static Source)

```
Interface:      xe26
Group:          239.1.1.1
Flags:          R
Uptime:         00:00:28
Group mode:     Include ()
State:          Active
Last reporter:  100.1.1.2
```

Group source list: (R - Remote, M - SSM Mapping, S - Static, L - Local)

Include Source List :

Source Address	Uptime	v3 Exp	Fwd	Flags
101.1.1.2	00:00:28	00:03:56	Yes	R

LHR#sh ipv6 mld groups

```
MLD Connected Group Membership
```

Group Address	Interface	Uptime	Expires	S
ff06::2	xe26	00:00:31	stopped	A

ctive fe80::1

LHR#**sh ipv6 mld groups detail**

MLD Connected Group Membership Details

Flags: (M - SSM Mapping, R - Remote,
SG - Static Group, SS - Static Source)

Interface: xe26

Group: ff06::2

Flags: R

Uptime: 00:00:32

Group mode: Include ()

State: Active

Last reporter: fe80::1

Group source list: (R - Remote, M - SSM Mapping, S - Static)

Include Source List :

Source Address	Uptime	v2 Exp	Fwd	Flags
5001::2	00:00:32	00:03:49	Yes	R

Improved Management

This section describes the network monitoring enhancements and new features introduced in the Release 6.5.1

- [Streaming Telemetry Dial-Out Mode](#)
- [DHCPv6 Prefix Delegation Configuration](#)
- [EVPN VXLAN E-Tree](#)

CHAPTER 1 Streaming Telemetry Dial-Out Mode

Overview

In OcNOS, dial-out telemetry subscriptions, also known as persistent subscriptions, ensure continuous data streaming, even if the Remote Procedure Call (gRPC) session terminates unexpectedly. With persistent subscriptions, the OcNOS device continuously retries to establish a gRPC connection to the collector server, thus maintaining persistent data streaming.

Feature Characteristics

The dial-out telemetry feature in OcNOS comprises several key aspects ensuring seamless data streaming and connectivity with collector servers:

The described topology outlines a system architecture that utilizes gRPC-based tunneling for persistent streaming telemetry.




Figure 1-6: Dial-Out Subscription Mode

Here is a detailed explanation of the components and data flow:

- **gNMI Client (gRPC Client):** The gNMI client, which acts as the gRPC client in this scenario, is responsible for handling telemetry data and connecting to the OcNOS target device.
- **Tunnel Server:** The tunnel server, part of the gNMI collector process, listens for incoming gRPC tunnel streams from the gRPC server.
- **gRPC Tunnel Stream:** Represents the persistent communication channel established between the tunnel client (OcNOS) and the tunnel server (collector).
- **Tunnel Client:** The gRPC tunnel client operates on the OcNOS device and connects to the tunnel server. It manages the tunneling of telemetry data.
- **gRPC Server:** Interacts with the tunnel client to establish and manage the tunnel.

Note: Ensure that the tunnel server is reachable over the network from the tunnel client, and configure both the tunnel client and tunnel server with compatible authentication mechanisms.

Data Flow

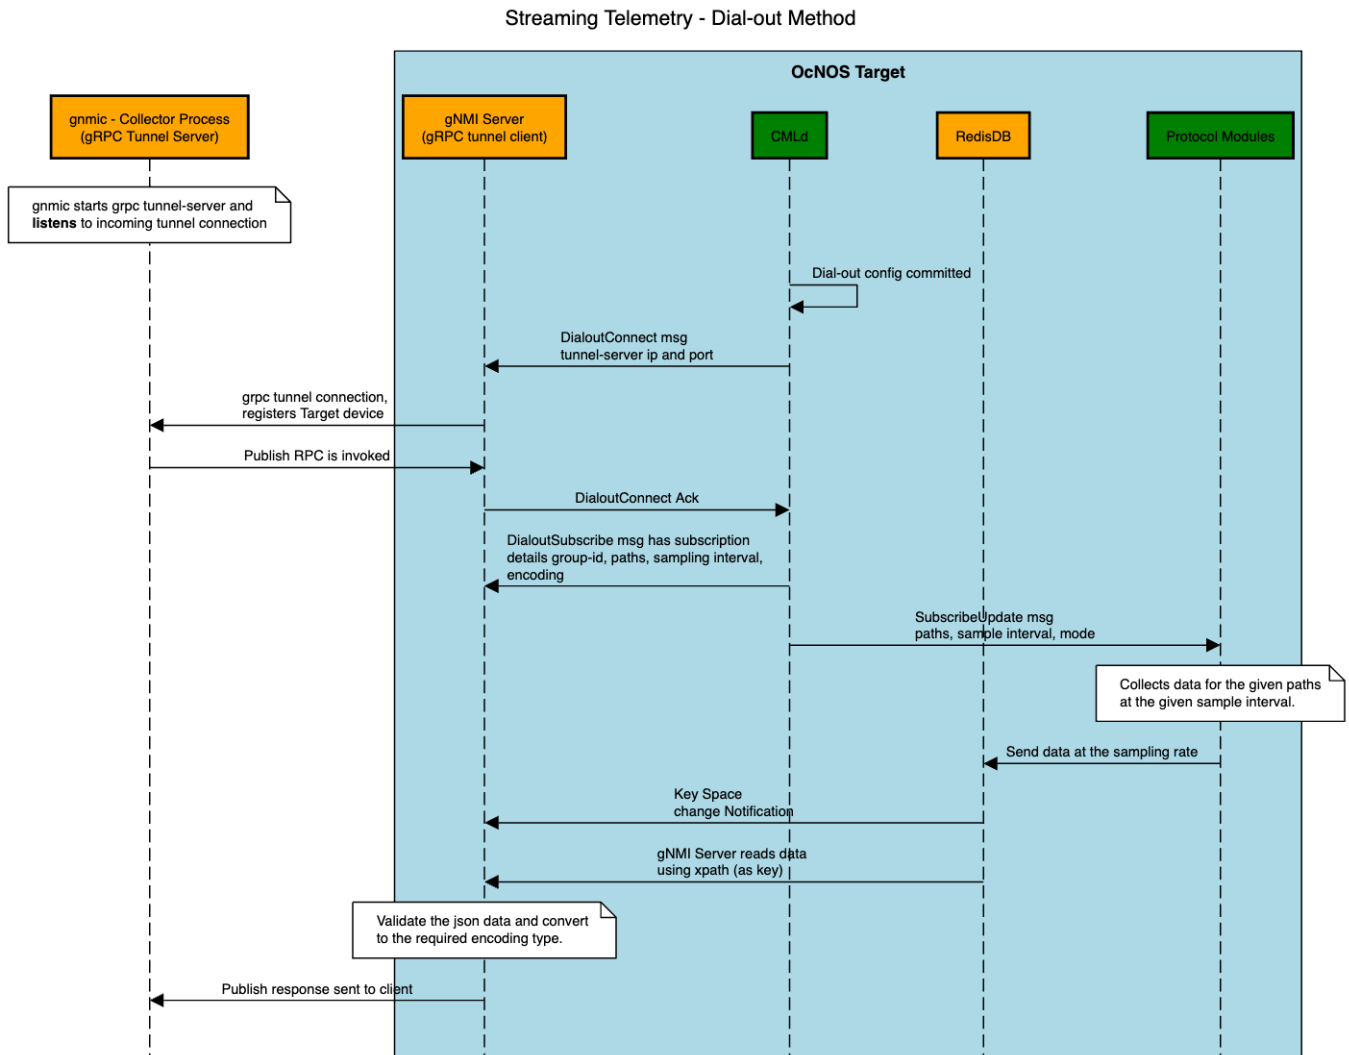
The [Data Flow: Dial-Out Mode](#) flow chart illustrates streaming telemetry in Dial-out Mode.

- **Initialization:** When the dial-out command `subscription-name` is applied successfully, the tunnel client on the OcNOS device initiates a connection to the tunnel server hosted on the collector.
- **Tunnel Establishment:** Upon successful connection, the gRPC client and server establish a persistent tunnel stream. This tunnel facilitates the continuous transmission of telemetry data.

Note: OcNOS supports insecure tunnel connections.

- **Telemetry Data Transmission:** When telemetry data needs to be transmitted from the OcNOS device, the gNMI client sends a Publish RPC request over the established tunnel.
- **Subscription Configuration:** Telemetry commands follow the OpenConfig telemetry model, standardizing the configuration of telemetry subscriptions and related entities.

Figure 1-7: Data Flow: Dial-Out Mode



Benefits

- Ensures continuous data streaming even in the event of gRPC session termination, enhancing network monitoring and troubleshooting capabilities.
- Simplifies configuration and management of telemetry subscriptions using standard OpenConfig models.
- Facilitates secure and reliable communication between the OcNOS device and the collector server.
- Enhances interoperability by enabling integration with third-party gRPC client applications like gNMI client, expanding telemetry options for network operators.

Prerequisites

Before configuring Dial-Out mode, ensure that:

- A supported OcnOS router running a compatible release is required.
- Access to the management interface of the router is necessary.
- Refer to the [gnmic Installation](#) to download the gNMI collector package.

Configuration

Set up the OcnOS router to transmit streaming telemetry data to a gNMI client using the dial-out method.

The sample configuration on the OcnOS router sets up streaming telemetry subscriptions using gNMI to monitor specific paths related to the state of Hard Disk and RAM. The router sends telemetry data to the specified collector (`Collector1`) over a configured tunnel connection. The gNMI client subscribed to these paths will receive updates regarding the state of RAM and Hard Disk at the specified intervals. This setup enables proactive monitoring and management of key hardware components on the network device.

Topology

In this setup, an OcnOS router functions as the data source for streaming telemetry, while a gNMI client acts as the receiver of telemetry data. The OcnOS router sends telemetry data to the gNMI client over a dial-out connection.



Figure 1-8: Dial-out Streaming Telemetry Topology

Configure OcnOS Router

Note: Before configuring Dial-out, meet all [Prerequisites](#).

1. Enable Streaming Telemetry on an OcnOS router.
`OcnOS(config)#feature streaming-telemetry`
2. Create Sensor Group

Create a sensor group (`Platform`) where sensor paths will be specified for dial-out subscriptions. Specify sensor paths within the sensor group (`Platform`) to monitor the state of RAM and Hard Disk.

```
OcnOS(config)#sensor-group Platform
OcnOS(telemetry-sensor-group)#sensor-path ipi:/components/component[name=RAM]/ram/state
```



```
OcNOS (telemetry-sensor-group) #sensor-path ipi:/components/component[name=HARD-DISK]/storage/state
OcNOS (telemetry-sensor-group) #exit
```

3. Create Destination Group

Create a destination group (`Collector1`) where tunnel server settings will be configured for dial-out subscriptions. Specify the tunnel server (`gnmi Client`) IP address (`10.12.101.72`) and port (`11161`) within the destination group (`Collector1`).

```
OcNOS (config) #destination-group Collector1
OcNOS (telemetry-grpc-tunnel-group) #tunnel-server ip 10.12.101.72 port 11161
OcNOS (telemetry-grpc-tunnel-group) #exit
```

4. Create Persistent Subscription

Create a persistent subscription (`storage`), encoding type (`JSON-IETF`), and associate it with the destination group (`Collector1`), and sensor group (`Platform`) to monitor the state of RAM and Hard Disk with a sample interval (`10` seconds).

```
OcNOS (config) #subscription-name storage
OcNOS (telemetry-subscription) #encoding json-ietf
OcNOS (telemetry-subscription) #destination-group Collector1
OcNOS (telemetry-subscription) #sensor-group Platform sample-interval 10
OcNOS (telemetry-subscription) #exit
```

Streaming Telemetry Snippet Configurations

To verify the telemetry configuration and view the overall commands used for dial-out subscriptions, use the `show running-config streaming-telemetry` command on the router.

```
OcNOS#show running-config streaming-telemetry
!
feature streaming-telemetry
debug telemetry gnmi enable severity debug
!
sensor-group storage
  sensor-path ipi:/components/component[name=RAM]/ram/state
  sensor-path ipi:/components/component[name=HARD-DISK]/storage/state
!
destination-group Collector1
  tunnel-server ip 10.12.101.72 port 11161
!
subscription-name State
  destination-group Collector1
  sensor-group storage sample-interval 10
!
!
```

Telemetry Subscription Invoked via `gnmic` Command and YAML Input

Start the `gnmic` collector with the `--use-tunnel-server` and `publish` options to receive the streamed gRPC responses. Execute the following command to start the gRPC tunnel server in listening mode, enabling it to accept incoming connections from gRPC tunnel clients (OcNOS target).

```
./gnmic --insecure --config <path to Tunnel-server yaml file> --use-tunnel-server publish
```

Invoke Publish RPC on OcNOS Target

The following output represents telemetry data published by the `gnmic` command, monitoring the state of Hard Disk and RAM on the specified OcNOS router.

```
# ./gnmic --insecure --config abc.yaml --use-tunnel-server publish
2024/04/12 11:22:50.516313 [gnmic] version=dev, commit=none, date=unknown,
gitURL=, docs=https://gnmic.openconfig.net
2024/04/12 11:22:50.516377 [gnmic] using config file "abc.yaml"
2024/04/12 11:22:50.517770 [gnmic] starting output type file
2024/04/12 11:22:50.517971 [file_output:default-stdout] initialized file
output:
{"Cfg":{"FileName":"","FileType":"stdout","Format":"json","Multiline":true,"In
dent":
","Separator":"\n","OverrideTimestamps":false,"AddTarget":"","TargetTemplate":
","EventProcessors":null,"MsgTemplate":"","ConcurrencyLimit":1000,"EnableMetric
s":false,"Debug":false}}
2024/04/12 11:22:50.518018 [gnmic] StartPublishCollector is invoked
2024/04/12 11:22:50.518446 [gnmic] Initializing error chan
2024/04/12 11:22:54.508410 [gnmic] tunnel server discovered target
{ID:e8:c5:7a:fe:fd:32 Type:GNMI_GNOI}
2024/04/12 11:22:54.508720 [gnmic] adding target
{"name":"e8:c5:7a:fe:fd:32","address":"e8:c5:7a:fe:fd:32","username":"root","p
assword":"****","timeout":1000000000,"insecure":true,"skip-
verify":false,"buffer-size":100,"retry-timer":1000000000,"log-tls-
secret":false,"gzip":false,"token":"","tunnel-target-type":"GNMI_GNOI"}
2024/04/12 11:22:54.508756 [gnmic] calling publishStream
2024/04/12 11:22:54.508772 [gnmic] publishStream is invoked
2024/04/12 11:22:54.508779 [gnmic] targetPublishStream is invoked
2024/04/12 11:22:54.508830 [gnmic] a.targetsChan: 0xc0004eb1a0
2024/04/12 11:22:54.508840 [gnmic] t.Config.Outputs: []
2024/04/12 11:22:54.508850 [gnmic] starting target "e8:c5:7a:fe:fd:32"
listener
2024/04/12 11:22:54.508879 [gnmic] queuing target "e8:c5:7a:fe:fd:32"
2024/04/12 11:22:54.508902 [gnmic] subscribing to target: "e8:c5:7a:fe:fd:32"
2024/04/12 11:22:54.508918 [gnmic] calling clientPublish
2024/04/12 11:22:54.508930 [gnmic] targetDialOpts: []grpc.DialOption
2024/04/12 11:22:54.508968 [gnmic] a.targetsChan: 0xc0004eb1a0
2024/04/12 11:22:54.508976 [gnmic] t.Config.Outputs: []
2024/04/12 11:22:54.509402 [gnmic] dialing tunnel connection for tunnel target
"e8:c5:7a:fe:fd:32"
Publish Request sent to e8:c5:7a:fe:fd:32{
  "source": "e8:c5:7a:fe:fd:32",
  "subscription-name": "storage",
  "timestamp": 1712920892603436151,
  "time": "2024-04-12T16:51:32.603436151+05:30",
  "updates": [
    {
      "Path": "ipi:components/component[name=HARD-DISK]/storage/state",
      "values": {
        "components/component/storage/state": {
          "free-memory": 0,
          "total-memory": 61057,
          "used-memory": 0
        }
      }
    }
  ]
}
{
```

```

"source": "e8:c5:7a:fe:fd:32",
"subscription-name": "storage",
"timestamp": 1712920892603253590,
"time": "2024-04-12T16:51:32.60325359+05:30",
"updates": [
  {
    "Path": "ipi:components/component[name=RAM]/ram/state",
    "values": {
      "components/component/ram/state": {
        "available-high-memory": 0,
        "available-memory": 15084,
        "buffers": 101,
        "current-process-count": 227,
        "free-swap": 0,
        "shared-memory": 28,
        "total-high-memory": 0,
        "total-memory": 16010,
        "total-swap": 0,
        "used-memory": 926
      }
    }
  }
]
}

```

The output of the Publish RPC includes the following information:

Publish RPC Output details

Option	Description
source	Displays the MAC address associated with the management port of the target. Each gNMI device have a unique target ID, allowing the collector to distinguish responses between various targets.
subscription-name	The name of the subscription.
timestamp	The timestamp of the response.
time	The timestamp in a human-readable format.
updates	An array of updates, each containing Path and Values.
Path	The path to the published data.
values	The values of the published data.

The telemetry data output includes detailed fields for monitoring the state of the Hard Disk and RAM, offering insights into the memory and storage utilization of the OcnOS router.

1. Hard Disk State

- **Free Memory:** The amount of free memory available on the hard disk.
- **Total Memory:** The total capacity of memory on the hard disk.
- **Used Memory:** The amount of memory currently in use on the hard disk.

2. RAM State

- **Available High Memory:** The available high memory in the RAM.

- **Available Memory:** The total available memory in the RAM.
- **Buffers:** The number of buffer processes running in the RAM.
- **Current Process Count:** The count of active processes in the RAM.
- **Free Swap:** The amount of free swap space in the RAM.
- **Shared Memory:** The shared memory usage in the RAM.
- **Total High Memory:** The total high memory capacity in the RAM.
- **Total Memory:** The total memory capacity in the RAM.
- **Total Swap:** The total swap space available in the RAM.
- **Used Memory:** The amount of memory currently in use in the RAM.

Validation

To verify persistent telemetry configurations and monitor the telemetry data transmission settings on the router, check the output of the `show streaming-telemetry persistent-subscriptions details` command.

```
#show streaming-telemetry persistent-subscriptions details
```

```
Feature streaming telemetry : Enabled
```

```
VRF : default
Platform type : High range
Maximum sensor-paths : 100
Minimum sample-interval : 10
Number of active sensor-paths : 2 (Dial-In : 0, Dial-out : 2)
Tunnel-server Retry-interval : Default-60 (seconds)
```

```
Enc-Type : Encoding type
SI : Sampling Interval in seconds
Origin:Path : Sensor Path
```

```
Dial-Out Subscription Details:
```

```
~~~~~
```

```
1. Subscription-name : State
   Status : ACTIVE
   Enc-Type : JSON
```

```
Tunnel-server details:
```

```
~~~~~
```

Destination-group	Status	Tunnel-IP:Port
Collector1	IN-ACTIVE	10.12.101.72:11161

```
Sensor-group details:
```

```
~~~~~
```

Sensor-group	SI	Origin:Path
storage	10	ipi:/components/component[name=RAM]/ram/state ipi:/components/component[name=HARD-DISK]/storage/state

Implementation Examples

Real-time Visibility: Operators have real-time visibility into network device health and performance metrics.

Proactive Maintenance: Early detection of issues allows for proactive maintenance and troubleshooting.

Optimized Resource Allocation: Insights from telemetry data help optimize resource allocation and capacity planning.

Enhanced Network Reliability: Continuous monitoring enhances network reliability and reduces downtime.

Dial-Out Commands

The streaming telemetry dial-out mode introduces the following configuration commands.

destination-group

Use this command to create a destination-group for persistent subscriptions on the OcNOS device. The VRF parameter must match the VRF specified in the [feature streaming-telemetry](#) command. Can create and attach multiple destination-groups to activate streaming telemetry subscriptions.

Use the no form of this command to delete a destination-group.

Command Syntax

```
destination-group TUNNEL-NAME (vrf (management|NAME) |)
no destination-group TUNNEL-NAME (vrf (management|NAME) |)
```

Parameters

TUNNEL-NAME	Specify the name assigned to the tunnel server or collector endpoint used for telemetry data transmission.
vrf NAME	(Optional) Creates a destination-group for persistent subscriptions in a user-defined VRF.
vrf	(Optional) Creates a destination-group for persistent subscriptions in the management VRF.
management	

Default

None

Command Mode

Configure Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example creates a destination group named `tunnel-1` for transmitting telemetry data.

```
OcNOS(config)#destination-group tunnel-1
OcNOS(telemetry-grpc-tunnel-group)#commit
```

destination-group GRPC

Use this command to add a destination-group under subscriptions. Can create multiple destination-groups within a subscription mode.

Use `no` parameter of this command to remove the destination-groups.

Note: Ensure that the GRPC-GROUP-NAME is configured in the device's configuration mode before adding it to a subscription mode.

Command Syntax

```
destination-group GRPC-GROUP-NAME
no destination-group GRPC-GROUP-NAME
```

Parameters

GRPC-GROUP-NAME	Specify the name assigned to the tunnel server or collector endpoint used for telemetry data transmission.
-----------------	--

Default

None

Command Mode

Telemetry-subscription Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

Ensure that the GRPC-GROUP-NAME (`tunnel-1`) is already configured in the current configuration mode.

```
OcNOS#configure terminal
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
grpc-tunnel-server retry-interval 60
!
sensor-group stream-1
  sensor-path ipi:/interfaces/interface[name=eth0]/state/counters
!
destination-group tunnel-1
  tunnel-server ip 10.12.66.160 port 11163
!
subscription-name sub-1
  sensor-group stream-1 sample-interval 1000
!
!
```

The following commands illustrates how to add a destination group (`tunnel-1`) under subscription mode (`sub-1`) and verify the configuration using the show command output.

```

OcNOS(config)#subscription-name sub-1
OcNOS (telemetry-subscription)#destination-group tunnel-1
OcNOS (telemetry-subscription)#commit
OcNOS (telemetry-subscription)#exit
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
grpc-tunnel-server retry-interval 60
!
sensor-group stream-1
  sensor-path ipi:/interfaces/interface[name=eth0]/state/counters
!
destination-group tunnel-1
  tunnel-server ip 10.12.66.160 port 11163
!
subscription-name sub-1
  destination-group tunnel-1
  sensor-group stream-1 sample-interval 1000
!
!

```

encoding

Use this command to specify or modify encoding types for subscriptions in streaming telemetry.

Use `no` parameter of this command to remove the encoding option.

Note: Modifying the encoding type is not allowed for active subscriptions.

Command Syntax

```

encoding (json-ietf|json|proto)
no encoding

```

Parameters

<code>json-ietf</code>	Specifies the JSON encoding based on the IETF draft standard.
<code>json</code>	Specifies the default JSON encoding type.
<code>proto</code>	Specifies the Protocol Buffers v3 encoding type.

Default

None

Command Mode

Telemetry-subscription Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following commands demonstrate how to create a telemetry subscription named `sub-3` using the JSON encoding type.

```
OcNOS#configure terminal
OcNOS(config)#subscription-name sub-3
OcNOS (telemetry-subscription)#encoding json
OcNOS (telemetry-subscription)#commit
```

grpc-tunnel-server retry-interval

Use this command to set the interval for retry attempts when establishing a connection for the GNMI server to the tunnel-server. The VRF parameter must match the VRF specified in the [feature streaming-telemetry](#) command.

Use `no` parameter of this command to unset the retry-interval timer.

Command Syntax

```
grpc-tunnel-server retry-interval <30-3000> (vrf (management|NAME) |)
no grpc-tunnel-server retry-interval (vrf (management|NAME) |)
```

Parameters

<code>retry-interval <30-3000></code>	Specifies the duration between retry attempts. The default retry-interval is 60 seconds.
<code>vrf management</code>	(Optional) Sets the retry-interval in the management VRF.
<code>vrf NAME</code>	(Optional) Sets the retry-interval in a user-defined VRF.

Default

None

Command Mode

Configure mode

Applicability

Introduced in the OcNOS version 6.5.2.

Example

The following configuration illustrates how to set the retry-interval timer for the gNMI server to the tunnel-server with a value of 80 seconds.

```
OcNOS#configure terminal
OcNOS(config)#feature streaming-telemetry
OcNOS(config)#grpc-tunnel-server retry-interval 80
OcNOS(config)#commit
```

sensor-group

Use this command to create a sensor group for persistent subscriptions in an OcNOS device. Multiple sensor groups can be created to specify the paths of interest for streaming telemetry. The VRF parameter must match the VRF specified in the [feature streaming-telemetry](#) command. These sensor groups are attached to subscriptions to activate streaming telemetry.

Use `no` parameter of this command to remove a created sensor group.

Command Syntax

```
sensor-group SENSOR-NAME (vrf (management|NAME) |)
no sensor-group SENSOR-NAME (vrf (management|NAME) |)
```

Parameters

<code>SENSOR-NAME</code>	Specifies the name of the sensor group.
<code>vrf</code>	(Optional) Creates a sensor group in the management VRF.
<code>management</code>	
<code>vrf NAME</code>	(Optional) Creates a sensor group in a user-defined VRF.

Default

None

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following commands demonstrate how to create a sensor group named "stream-1" for persistent telemetry subscriptions on an OcNOS device:

```
OcNOS#configure terminal
OcNOS(config)#sensor-group stream-1
OcNOS(telemetry-sensor-group)#commit
OcNOS(telemetry-sensor-group)#exit
```

sensor-group sample-interval

Use this command to to associate a sensor group with a specific sampling interval under subscriptions for activating streaming telemetry. Multiple sensor groups can be created.

Use `no` parameter of this command to remove the sensor-groups from a subscription.

Note: Before adding a `SENSOR-GROUP-NAME` to a subscription, ensure the sensor group is already configured in the configuration mode.

Command Syntax

```
sensor-group SENSOR-GROUP-NAME sample-interval <10-3600>
no sensor-group SENSOR-GROUP-NAME
```

Parameters

SENSOR-GROUP-NAME	Specifies the name of the sensor group to be associated with the subscription.
sample-interval <10-3600>	Defines the sampling interval in seconds for the sensor group. The interval can range from 10 to 3600 seconds.

Default

None

Command Mode

Telemetry-subscription Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

Ensure that the SENSOR-GROUP-NAME (`stream-1`) is already configured in the current configuration mode.

```
OcNOS#configure terminal
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
grpc-tunnel-server retry-interval 60
!
sensor-group stream-1
  sensor-path ipi:/interfaces/interface[name=eth0]/state/counters
!
subscription-name sub-1
!
!
```

The following commands illustrates how to add a sensor group (`stream-1`) under subscription mode (`sub-1`) and verify the configuration using the show command output.

```
OcNOS(config)#subscription-name sub-1
OcNOS(telemetry-subscription)#sensor-group stream-1 sample-interval 1000
OcNOS(telemetry-subscription)#commit
OcNOS(telemetry-subscription)#exit
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
grpc-tunnel-server retry-interval 60
!
sensor-group stream-1
  sensor-path ipi:/interfaces/interface[name=eth0]/state/counters
!
subscription-name sub-1
sensor-group stream-1 sample-interval 1000
!
!
```

sensor-path

Use this command to add sensor paths under sensor-groups. Can add multiple sensor paths to a single sensor group.

Use `no` parameter of this command to remove sensor paths.

Command Syntax

```
sensor-path SENSOR-PATH
no sensor-path SENSOR-PATH
```

Parameters

`SENSOR-PATH` Specifies the path of the telemetry data to include in the sensor group.

Default

None

Command Mode

Telemetry-sensor-group Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following example demonstrates how to configure a sensor group (`stream-1`) and add multiple sensor paths to it for streaming telemetry.

```
OcNOS#configure terminal
OcNOS(config)#sensor-group stream-1
OcNOS(telemetry-sensor-group)#sensor-path ipi:/interfaces/
interface[name=eth0]/state/counters
OcNOS(telemetry-sensor-group)#sensor-path /interfaces/interface[name=xe2]/
state/counters
OcNOS(telemetry-sensor-group)#sensor-path openconfig:/interfaces/
interface[name=xe3]/state/counters
OcNOS(telemetry-sensor-group)#commit
OcNOS(telemetry-sensor-group)#exit
```

show streaming-telemetry persistent-subscriptions

Use this command to display a brief summary of the streaming-telemetry dial-out configurations. This command provides a concise view of the persistent subscription settings configured on the device.

Command Syntax

```
show streaming-telemetry persistent-subscriptions brief
show streaming-telemetry persistent-subscriptions details (SUBSCRIPTION-NAME|)
```

Parameters

SUBSCRIPTION-NAME Displays detailed configuration information specific to the named persistent subscription.

Default

None

Command Mode

Exec mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The command output lists each persistent subscription with its associated details.

```
OcNOS#show streaming-telemetry persistent-subscriptions details
```

```
Feature streaming telemetry : Enabled
```

```
VRF : default
Platform type : High range
Maximum sensor-paths : 100
Minimum sample-interval : 10
Number of active sensor-paths : 2 (Dial-In : 0, Dial-out : 2)
Tunnel-server Retry-interval : Default-60 (seconds)
```

```
Enc-Type : Encoding type
SI : Sampling Interval in seconds
Origin:Path : Sensor Path
```

```
Dial-Out Subscription Details:
```

```
~~~~~
```

```
1. Subscription-name : State
   Status : ACTIVE
   Enc-Type : JSON
   Tunnel-server details:
```

```
~~~~~
```

Destination-group	Status	Tunnel-IP:Port
Collector1	IN-ACTIVE	10.12.101.72:11161

```
Sensor-group details:
```

```
~~~~~
```

Sensor-group	SI	Origin:Path
storage	10	ipi:/components/component [name=RAM] /ram/state ipi:/components/component [name=HARD-DISK] /storage/state

The following table explains the output fields.

Field	Description
Feature streaming telemetry	Marked as “Enabled” confirms that streaming telemetry is active on the device.
VRF	Specifies the VRF type.
Platform type	Displays the platform type is standard or high range.
Maximum sensor-paths	Shows the maximum number of sensor paths allowed. For more details, refer to Scale Scenarios section.
Minimum sample-interval	Indicates the minimum sampling interval in seconds. For more details, refer to Scale Scenarios section.
Number of active sensor-paths	Shows the total number of active sensor paths for Dial-In and Dial-Out subscriptions (Stream mode subscriptions).
Tunnel-server Default-Retry-interval	The duration between retry attempts when establishing a connection for the GNMI server to the tunnel server.
Subscription Name	Name of the persistent subscription.
Storage Status or Status	Current status of the subscription (ACTIVE or IN-ACTIVE).
Enc-Type	Encoding type used for telemetry data (JSON, JSON-IETF, Proto).
Destination Group	Define the tunnel server settings to which telemetry data is sent for dial-out subscriptions.
Sensor Group	Sensor group associated with the subscription.
Sample Interval (SI)	Sampling interval for the sensor group.
Tunnel-IP:Port	IP address and port of the tunnel server for dial-out subscriptions.
Origin:Path	The specific sensor paths that are being monitored or streamed by the telemetry system.

subscription-name

Use this command to create named subscriptions for persistent telemetry configurations in an OcNOS device. The VRF parameter must match the VRF specified in the [feature streaming-telemetry](#) command. Multiple subscriptions can be created. These subscriptions are essential for activating streaming telemetry, as they define specific settings such as associated destination groups and sensor groups.

Use `no` parameter of this command to delete a subscription.

Command Syntax

```
subscription-name NAME (vrf (management|NAME) |)
no subscription-name NAME (vrf (management|NAME) |)
```

Parameters

subscription-name NAME	Specifies the unique name to the persistent subscription.
vrf NAME	(Optional) Creates named subscriptions in a user-defined VRF.
vrf management	(Optional) Creates named subscriptions in the management VRF.

Default

None

Command Mode

Configure Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following command demonstrates configuring a subscription (`sub-1`) on an OcNOS device. The subscription remains `in-active` because the sensor groups and destination groups have not been added to it.

```
OcNOS#configure terminal
OcNOS(config)#subscription-name sub-1
OcNOS(telemetry-subscription)#commit
Subscription sub-1 is "in-active": sensor-group(s) and destination-group(s)
are not configured.
OcNOS(telemetry-subscription)#exit
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
!
subscription-name sub-1
!
!
```

tunnel-server

Use this command to add tunnel-servers under destination groups. Can create multiple tunnel servers within a destination group.

Use `no` parameter of this command to remove a tunnel server from the destination group.

Command Syntax

```
tunnel-server ip A.B.C.D port <1-65535>
no tunnel-server ip A.B.C.D port <1-65535>
```

Parameters

<code>ip A.B.C.D</code>	Specifies the tunnel server IP address.
<code>port <1-65535></code>	Specifies the tunnel server port-number.

Default

None

Command Mode

Telemetry-GRPC-tunnel-group Mode

Applicability

Introduced in OcNOS version 6.5.2.

Example

The following command demonstrates how to add a tunnel server within the destination group.

```
OcNOS#configure terminal
OcNOS(config)#destination-group tunnel-1
OcNOS(telemetry-grpc-tunnel-group)#tunnel-server ip 10.12.66.160 port 11163
OcNOS(telemetry-grpc-tunnel-group)#commit
OcNOS(telemetry-grpc-tunnel-group)#exit
```

Revised CLI Commands

The following is the revised command for telemetry.

show techsupport

- The existing syntax now includes the newly added parameter for telemetry, namely `gnmi`.
- The command `show techsupport gnmi` collects gNMI-related information for technical support. For more details, refer to the `show techsupport` command in the Software Monitoring and Reporting chapter in the *System Management Guide*.

Glossary

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

Key Terms/Acronym	Description
Remote Procedure Call (gRPC)	gRPC protocol that uses HTTP/2 for transport and protocol buffers for serialization.
Persistent Subscription	Telemetry subscription that maintains continuous data streaming even after interruptions in connectivity.

gRPC Network Management Interface (gNMI)	A standardized protocol for network management using gRPC and protocol buffers.
Destination Group	Specifies the collector server's details and connection parameters for telemetry subscriptions.
Sensor Group	Contains sensor paths that define the specific data to be monitored and transmitted.
OpenConfig	Standardized model for network configuration and telemetry using a vendor-neutral approach.

CHAPTER 2 DHCPv6 Prefix Delegation Configuration

Overview

The prefix delegation feature facilitates the Dynamic Host Control Protocol (DHCP) server capable of assigning prefixes to DHCP clients from a global pool, enabling the Customer Premise Equipment (CPE) to learn the prefix. This feature also supports the DHCP server in assigning multiple prefixes to a single client. The user configures the IPv6 address using the learned prefix on its Local Area Network (LAN) interface with the subnet prefix. The LAN hosts are learning the subnetted prefix through Router Advertisement (RA) messages, an important Neighbor Discovery Protocol (NDP) component, enabling the device to auto-configure the number of IPv6 addresses from 1 to 64.

This feature would enable service providers to assign IP for the CPE that is acting as a router between the service providers' core network and the subscribers' internal network.

Feature Characteristics

- DHCPv6 Identity association for non-temporary addresses (IA_NA) assigns a global IPv6 address on the Wide Area Network (WAN) link. The address comes from a local pool specified in the DHCP Server.
- The Requesting Router (RR) uses the delegated prefix to define the subnet for the LAN based on the prefix received from the DHCP Server.
- The Requesting Router uses the delegated prefix to assign addresses to the LAN devices. The RR can send a Router Advertisement or the devices shall send a Router solicitation.

Benefits

The key benefits are as follows:

- This feature helps the Internet Service Providers (ISPs) to assign the dynamic IPv6 addresses to their customers automatically instead of statically assigning the address.
- This feature adds the capability to get the multiple DHCPv6 prefixes as per the customer requirement.
- This feature allows the centralized management of the IPv6 addresses.

Configuration

This section shows the configuration of the DHCPv6 prefix delegation.

Topology

The requesting router sends the prefix request to the delegating router, which sends the request to the DHCP server. The DHCP server sends the prefix to the requesting router through the delegating router. The IPv6 address is created in the requesting router by combining the prefix learned from the server and the user-defined suffix. The host receives the IPv6 address from the requesting router.




Figure 2-9: DHCPv6 Prefix Delegation Configuration

Configuring DHCP prefixes

Follow the steps to configure the DHCPv6 prefix delegation.

Configure the Delegating Router:

1. Specify the server interface address connected to the delegating router.


```
(config)#ipv6 dhcp relay address 2001:101:0:1::131
```
2. Configure the DHCPv6 up-link interface from the delegating router to the DHCPv6 server using `ipv6 dhcp relay uplink` command.


```
(config)#interface ce1/2
(config-if)#ipv6 address 2001:101:0:1::130/64
(config-if)#ipv6 dhcp relay uplink
```
3. Configure the DHCPv6 down-link interface from the delegating router to the requesting router using `ipv6 dhcp relay` command.


```
(config)#interface ce16/1
(config-if)#ipv6 address 3001:101:0:1::135/64
(config-if)#ipv6 dhcp relay
```
4. Add a static route on the delegating router to reach the host device.


```
(config)#ipv6 route ::/0 3001:101:0:1::
```

Configure the Requesting Router device:

1. In the WAN interface, configure the address prefix length option (64). Get the IPv6 address from the server using `ipv6 address dhcp` command. Enable the requesting router to request the prefix by using `ipv6 dhcp prefix-delegation` and configure the number of prefixes using `ipv6 dhcp client max-delegated-prefixes`.

Note: The default value of simultaneous prefixes delegated to a single client is 8. The minimum of simultaneous prefixes delegated to a single client is 1 and the maximum is 64.

Note: If the configured `max-delegated-prefix count` is greater than 30, then configure the lease times greater than 180 seconds.

```
(config)#interface ce16/1
(config-if)#ipv6 dhcp address-prefix-len 64
(config-if)#ipv6 address dhcp
(config-if)#ipv6 dhcp prefix-delegation PREFIX_FROM_SERVER
(config-if)#ipv6 dhcp client max-delegated-prefixes 10
```

2. In the LAN interface, configure the command `ipv6 address` to create the IPv6 address by using the DHCP prefix learned from the server and user defined suffix.

```
(config)#interface ce23/1
(config-if)#ipv6 address PREFIX_FROM_SERVER ::1:0:0:0:1/64
```

3. Add a static route on the requesting router to reach the host device.

```
(config)#ipv6 route 2001:101:0:1::/64 3001:101:0:1::135
```

Configure the HOST:

1. In the LAN interface, configure the auto-configuration to get the dynamic IPv6 address from the server.

```
(config)#interface eth3
(config-if)#ipv6 address autoconfig max-address 10
(config if)#exit
(config)#commit
```

2. Add a static route on the host to reach the server.

```
(config)#ipv6 route 2001:101:0:1::/64 3001:101:0:1::135
```

Running configurations

The running configuration for the Delegating Router is as follows:

```
#show running-config
!
ipv6 dhcp relay address 2001:101:0:1::131
!
interface ce1/2
  ipv6 address 2001:101:0:1::130/64
  ipv6 dhcp relay uplink
!
interface ce16/1
  ipv6 address 3001:101:0:1::135/64
  ipv6 dhcp relay
  commit
end
!
```

The running configuration for the Requesting Router is as follows:

```
#show running-config
!
interface ce16/1
  ipv6 dhcp client max-delegated-prefixes 10
  ipv6 address dhcp
  ipv6 dhcp address-prefix-len 64
  ipv6 dhcp prefix-delegation PREFIX_FROM_SERVER
!
interface ce23/1
  ipv6 address PREFIX_FROM_SERVER ::1:0:0:0:1/64
  commit
end
!
```

The running configuration for the HOST is as follows:

```
#show running-config
!
interface eth3
```

```

    ipv6 address autoconfig max-address 10
    commit
end
!
```

Validation

Validate the show output after configuration as shown below.

Delegating Router:

```

#show ipv6 route
IPv6 Routing Table
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       v - vrf leaked
Timers: Uptime

IP Route Table for VRF "default"
C       ::1/128 via ::, lo, 00:03:20
C       2001:101:0:1::/64 via ::, ce16/2, 00:02:58
D       2001:db9:c0f::/48 [80/0] via fe80::eac5:7aff:fe51:723b, ce16/1, 00:00:44
C       3001:101:0:1::/64 via ::, ce16/1, 00:00:50
C       fe80::/64 via ::, ce16/1, 00:00:50
#show ipv6 dhcp pd-route
VRF : default
  2001:db9:c0a::/48 via 2001:db9:c0b::, ce16/1, (2024-03-07 06:20:43 - 2024-03-07
06:22:13)
  2001:db9:c0b::/48 via 2001:db9:c09::, ce16/1, (2024-03-07 06:20:42 - 2024-03-07
06:22:12)
  2001:db9:c0c::/48 via 2001:db9:c0d::, ce16/1, (2024-03-07 06:20:39 - 2024-03-07
06:22:09)
  2001:db9:c0d::/48 via 2001:db9:c0e::, ce16/1, (2024-03-07 06:20:38 - 2024-03-07
06:22:08)
  2001:db9:c0e::/48 via 2001:db9:c0f::, ce16/1, (2024-03-07 06:20:37 - 2024-03-07
06:22:07)
  2001:db9:c0f::/48 via fe80::eac5:7aff:fe51:723b, ce16/1, (2024-03-07 06:20:36 - 2024-
03-07 06:22:06)
  2001:db9:c05::/48 via 2001:db9:c06::, ce16/1, (2024-03-07 06:20:45 - 2024-03-07
06:22:15)
  2001:db9:c06::/48 via 2001:db9:c0a::, ce16/1, (2024-03-07 06:20:44 - 2024-03-07
06:22:14)
  2001:db9:c08::/48 via 2001:db9:c0c::, ce16/1, (2024-03-07 06:20:40 - 2024-03-07
06:22:10)
  2001:db9:c09::/48 via 2001:db9:c08::, ce16/1, (2024-03-07 06:20:41 - 2024-03-07
06:22:11)
#show ipv6 dhcp relay
IPv6 DHCP relay service is Enabled.
VRF Name: default
  DHCPv6 Servers configured:
    2001:101:0:1::131
```

```
DHCPv6 IA_PD Route injection: Enabled
DHCPv6 Duplicate Clients detection: Disabled
Interface                Uplink/Downlink
-----                -
ce16/1                   Downlink
ce1/2                    Uplink
```

Requesting Router:

```
#show ipv6 dhcp interface
```

```
ce16/1 is in client mode
  prefix name: PREFIX_FROM_SERVER
  learned prefix: 2001:db9:c05::/48
  preferred lifetime 0, valid lifetime 60
  interfaces using the learned prefix
    ce23/1    2001:db9:c0f:1::1
    ce23/1    2001:db9:c0e:1::1
    ce23/1    2001:db9:c0d:1::1
    ce23/1    2001:db9:c0c:1::1
    ce23/1    2001:db9:c08:1::1
    ce23/1    2001:db9:c09:1::1
    ce23/1    2001:db9:c0b:1::1
    ce23/1    2001:db9:c0a:1::1
    ce23/1    2001:db9:c06:1::1
    ce23/1    2001:db9:c05:1::1
```

```
#show interface ce23/1
```

```
Interface ce23/1
  Flexport: Non Control Port (Active)
  Hardware is ETH Current HW addr: e8c5.7a51.722e
  Physical:e8c5.7a51.722e Logical:(not set)
  Forward Error Correction (FEC) configured is Auto (default)
  FEC status is N/A
  Port Mode is Router
  Protected Mode is Promiscuous
  Interface index: 10017
  Metric 1 mtu 1500 duplex-full link-speed 10g
  Debounce timer: disable
  ARP ageing timeout 1500
  <UP,BROADCAST,RUNNING,ALLMULTI,MULTICAST>
  VRF Binding: Not bound
  Label switching is disabled
  No Virtual Circuit configured
  Administrative Group(s): None
  Bandwidth 10g
  Maximum reservable bandwidth 10g
    Available b/w at priority 0 is 10g
    Available b/w at priority 1 is 10g
    Available b/w at priority 2 is 10g
    Available b/w at priority 3 is 10g
```

```

    Available b/w at priority 4 is 10g
    Available b/w at priority 5 is 10g
    Available b/w at priority 6 is 10g
    Available b/w at priority 7 is 10g
DHCP client is disabled.
Last Flapped: Never
Statistics last cleared: Never
inet6 2001:db9:c05:1::1/64
inet6 2001:db9:c06:1::1/64
inet6 2001:db9:c08:1::1/64
inet6 2001:db9:c09:1::1/64
inet6 2001:db9:c0a:1::1/64
inet6 2001:db9:c0b:1::1/64
inet6 2001:db9:c0c:1::1/64
inet6 2001:db9:c0d:1::1/64
inet6 2001:db9:c0e:1::1/64
inet6 2001:db9:c0f:1::1/64
inet6 fe80::eac5:7aff:fe51:722e/64
ND router advertisements are sent approximately every 561 seconds
ND next router advertisement due in 517 seconds.
ND router advertisements live for 1800 seconds
Hosts use stateless autoconfig for addresses.
5 minute input rate 82 bits/sec, 0 packets/sec
5 minute output rate 191 bits/sec, 0 packets/sec
RX
    unicast packets 0 multicast packets 25 broadcast packets 0
    input packets 25 bytes 2862
    jumbo packets 0
    undersize 0 oversize 0 CRC 0 fragments 0 jabbers 0
    input error 0
    input with dribble 0 input discard 0
    Rx pause 0
TX
    unicast packets 0 multicast packets 38 broadcast packets 0
    output packets 38 bytes 5540
    jumbo packets 0
    output errors 0 collision 0 deferred 0 late collision 0
    output discard 0
    Tx pause 0

```

HOST:

```

#show ipv6 interface eth3 brief
Interface                IPv6-Address                Admin-Status
eth3                      2001:db9:c05:1:923c:b3ff:fe90:9fa9
                          2001:db9:c06:1:923c:b3ff:fe90:9fa9
                          2001:db9:c08:1:923c:b3ff:fe90:9fa9
                          2001:db9:c09:1:923c:b3ff:fe90:9fa9
                          2001:db9:c0a:1:923c:b3ff:fe90:9fa9
                          2001:db9:c0b:1:923c:b3ff:fe90:9fa9
                          2001:db9:c0c:1:923c:b3ff:fe90:9fa9

```

```

2001:db9:c0d:1:923c:b3ff:fe90:9fa9
2001:db9:c0e:1:923c:b3ff:fe90:9fa9
2001:db9:c0f:1:923c:b3ff:fe90:9fa9
fe80::923c:b3ff:fe90:9fa9

```

[up/up]

DHCP Multiple Prefix Delegation Command

The DHCPv6 Prefix Delegation introduces the following configuration command.

ipv6 dhcp client max-delegated-prefixes

Use this command to configure multiple DHCPv6 prefix delegation for a single client.

Command Syntax

```
ipv6 dhcp client max-delegated-prefixes <1-64>
```

Parameters

<pre>max- delegated- prefixes <1- 64></pre>	<p>Specifies the number of prefixes need for a DHCP client. Default number of DHCP prefixes are 8.</p>
---	--

Default

None

Command Mode

Interface mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

This example shows how to configure multiple DHCPv6 prefix delegation for a single client:

```

RR#configure terminal
RR#(config)#interface ce16/1
RR#(config-if)#ipv6 dhcp address-prefix-len 64
RR#(config-if)#ipv6 address dhcp
RR#(config-if)#ipv6 dhcp prefix-delegation PREFIX_FROM_SERVER
RR#(config-if)#ipv6 dhcp client max-delegated-prefixes 10
RR#(config-if)#exit
RR#(config)#commit

```

Revised CLI Commands

The following command is revised:

ipv6 address autoconfig

The existing syntax now includes the newly added parameter (`max-address <1-64>|`). For more details, refer to [ipv6 address autoconfig](#) command in the [DHCPv6 Prefix delegation Commands](#) chapter in the *System Management Guide*.

Glossary

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

Key Terms/Acronym	Description
Border Network Gateway (BNG)	Border Network Gateway is a critical component in the telecommunication network that serves as the entry and exit point between the ISP and the global network.
Customer Premises Equipment (CPE)	Customer Premises Equipment is a networking device located on the customer premises. It is present on the edge of the service provider network, which connects the customer devices to the service provider network.
Delegating Router (DR)	Delegating Router is a network device that delegates the IPv6 address prefixes to the downstream devices.
Identity association for non-temporary addresses (IA_NA)	Identity association for non-temporary addresses is a unique identifier associated with a set of IPv6 addresses assigned to client devices permanently or for a long time.
Local Area Network (LAN)	Local Area Network is a network of devices in a small area that may include a building or home.
Neighbor Discovery Protocol (NDP)	Neighbor Discovery Protocol is a crucial protocol in the IPv6 networks, helping establish the communication and auto-configuration to run the devices in the local network segment seamlessly.
Neighbor Discovery Router Advertisement (NDRA)	Neighbor Discovery Router Advertisement facilitates a network device to advertise the routing information with the neighboring devices so that the neighboring devices take the forwarding decision in dynamic routing.
Router Advertisement (RA)	Router Advertisement is a critical component in the IPv6 network. The router sends a message to the devices connected to the LAN to communicate its presence and share the configurations with the LAN host.
Requesting Router (RR)	Requesting Router is a network device that requests the IPv6 address prefixes to the DHCP server to share it with the downstream devices.
Router Solicitation (RS)	Router Solicitation is a component of the neighbor discovery protocol in the IPv6 network where the host sends a message to discover routers in the local area. When a router receives RS, it responds to the host with RA, which includes the configuration.
Wide Area Network (WAN)	Wide Area Network refers to large network that includes multiple LANs and spans over a large geographical area.

CHAPTER 3 EVPN VXLAN E-Tree

Overview

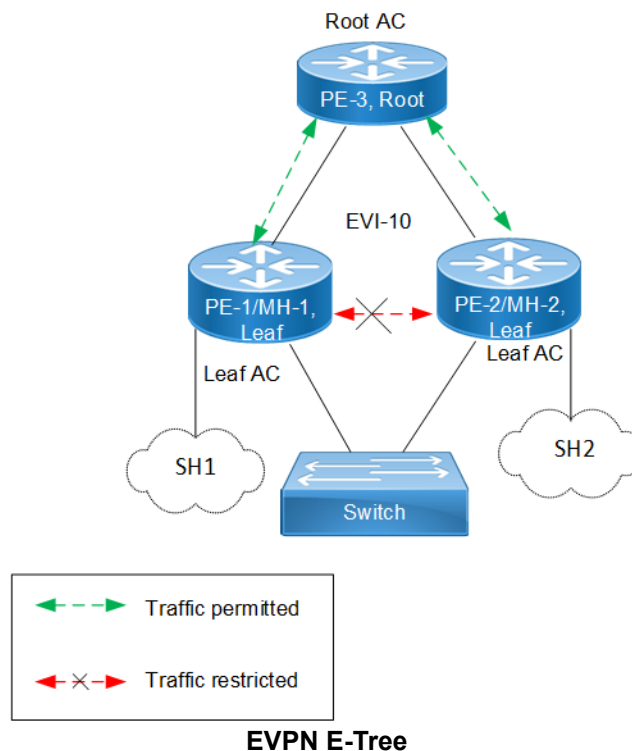
Ethernet VPN Ethernet-Tree (EVPN E-Tree), is a networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, especially in scenarios where specific services or devices need controlled communication. It categorizes network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them.

Feature Characteristics

Implemented Scenario 1 of the EVPN E-Tree solution, as defined by RFC-8317, designates each Provider Edge (PE) node as either a Leaf or a Root site per Virtual Private Network (VPN) for VXLAN EVPN in OcnOS.

Scenario 1: Leaf or Root Site(s) per PE

The explanation of scenario 1 is based on the provided topology diagram, which consists of three PE nodes labeled PE-1, PE-2, and PE-3 and two Multi-Homed (MH) nodes labeled MH-1 and MH-2. Within this setup, PE-3 functions as the Root node, while PE-1 and PE-2 serve as Leaf nodes. Also, PE-1 and PE-2 are part of a single home access-if port (SH1 and SH2).



The classification ensures that communication follows specific rules:

- Communication between Leaf hosts is restricted, as indicated by red dotted lines with a cross mark (X) in the topology diagram. However, communication between Leaf and Root nodes, as well as between Root nodes, is permitted, marked by green dotted lines.

- Leaf nodes within PE-1 and PE-2 are isolated from each other, preventing intra-PE communication.

The scenario 1 is achieved through two main concepts:

1. Inter-PE Communication

- The inter-PE Route Target (RT) Constraint Method is applicable only to Single-Homing (SH) devices. Two RTs per broadcast domain are utilized, with Leaf PEs exporting Leaf RTs and Root nodes exporting Root RTs. Leaf nodes import only Root RTs, allowing communication with Root PEs while preventing communication with other Leaf nodes. RT constraints limit the import of specific EVPN routes (MAC-IP and IMET routes) to designated paths for inter-PE communication.
- IPI employs a proprietary method to support inter-PE connectivity for both SH and MH devices, using BGP extended community to advertise Leaf Indication in BGP routes and influence traffic flow for both Unicast and BUM traffic. This method enables implementation of ARP or ND cache suppression and MAC mobility sub-features specified in RFC-7432.

2. Intra-PE communication: Local Split Horizon controls intra-PE communication between Attachment Circuits (ACs) within Leaf PE nodes, ensuring that traffic between ACs does not egress to other Leaf ACs.

Note: This functionality depends on hardware capabilities.

Benefits

EVPN E-Tree offers benefits in networking environments by providing efficient traffic control, enhanced security, scalability, and improved performance.

Efficient Traffic Control: EVPN E-Tree allows for efficient control over traffic within network broadcast domains. By segregating nodes into Leaf and Root categories, it enables precise management of communication flows, ensuring the traffic is directed only where needed.

Enhanced Security: The isolation of Leaf hosts from each other adds a layer of security to the network. This prevents unauthorized communication between devices within the same broadcast domain, reducing the risk of data breaches and unauthorized access.

Scalability: EVPN E-Tree is scalable, making it suitable for networks of various sizes and complexities. Whether deploying in small-scale environments or large enterprise networks, EVPN E-Tree offers flexibility and scalability to meet evolving business needs.

Improved Performance: By controlling communication paths and optimizing traffic flows, EVPN E-Tree can improve network performance. This ensures that critical data packets are delivered efficiently, reducing latency and enhancing overall network performance.

Prerequisites

In setting up a VXLAN EVPN network, certain prerequisites are essential to ensure proper functionality and connectivity.

Ensure VXLAN EVPN Configuration: Confirm that VXLAN, EVPN VXLAN, and VXLAN filtering are already enabled in the network as they are required for VXLAN EVPN Multihoming.

Define Interfaces and Loopback Addresses: Configure Layer 2 interfaces, like port channel interfaces (e.g., po1), and assign specific system MAC addresses (Ethernet Segment Identifier (ESI) values) for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity. These configurations establish the efficient network routing and communication.

Configure OSPF and BGP for Dynamic Routing: Enable OSPF to facilitate dynamic routing within the network. Define OSPF router IDs to match loopback IP addresses and add network segments to OSPF areas for proper route

distribution. Additionally, establish BGP sessions to advertise routes between different nodes. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

Leaf Node

1. Enable VXLAN and EVPN MH

Enable features like VXLAN and EVPN Multihoming, VXLAN filtering, and quality of service (QoS) capabilities on all Leaf nodes.

```
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
nvo vxlan enable
!
evpn vxlan multihoming enable
!
qos enable
!
```

2. Configure Interfaces and Loopback

Define a port channel interface (po1) as an L2 interface and assign the system MAC (0000.0000.1111) as the ESI value. Designate an interface (xe7) as a member port of po1. Assign the loopback IP address (1.1.1.1) to Leaf node, and set IP addresses (10.10.10.1 and 10.10.11.1) to interfaces (xe45 and xe49/2), respectively, for connectivity with Spine nodes.

```
!
interface po1
  switchport
  evpn multi-homed system-mac 0000.0000.1111
!
interface lo
  ip address 1.1.1.1/32 secondary
!
interface xe7
  channel-group 1 mode active
!
interface xe45
  ip address 10.10.10.1/24
!
interface xe49/2
  ip address 10.10.11.1/24
  exit
!
```

3. Configure OSPF

In OSPF router mode, set the router ID (1.1.1.1), to match the loopback IP address. Add the loopback network (1.1.1.1/32) and networks (10.10.10.0/24 and 10.10.11.0/24) connected to Spine nodes in OSPF area 0. Enable Bidirectional Forwarding Detection (BFD) on all OSPF interfaces for faster convergence.

```
!
router ospf 100
  ospf router-id 1.1.1.1
  bfd all-interfaces
  network 1.1.1.1/32 area 0.0.0.0
```

```

network 10.10.10.0/24 area 0.0.0.0
network 10.10.11.0/24 area 0.0.0.0
!
```

4. Configure BGP

In BGP router mode, set the router ID (1.1.1.1) to match the loopback IP address. Specify the loopback IP address of each Leaf node as neighbors with their respective remote AS numbers. Configure the loopback as the update source for each neighbor and set the advertisement interval (0) for rapid convergence. In L2VPN EVPN address family mode, activate each Leaf node (2.2.2.2, 3.3.3.3, 4.4.4.4) to establish connections within the EVPN address family.

```

!
router bgp 100
  bgp router-id 1.1.1.1
  neighbor 2.2.2.2 remote-as 100
  neighbor 3.3.3.3 remote-as 100
  neighbor 4.4.4.4 remote-as 100
  neighbor 2.2.2.2 update-source lo
  neighbor 2.2.2.2 advertisement-interval 0
  neighbor 3.3.3.3 update-source lo
  neighbor 3.3.3.3 advertisement-interval 0
  neighbor 4.4.4.4 update-source lo
  neighbor 4.4.4.4 advertisement-interval 0
!
address-family l2vpn evpn
  neighbor 2.2.2.2 activate
  neighbor 3.3.3.3 activate
  neighbor 4.4.4.4 activate
exit-address-family
!
exit
!
```

5. Configure VRF

In VRF mode, create a MAC routing or forwarding instance (VRF1). Assign the Route Distinguisher (RD) value (1.1.1.1:100) and set both import and export route-target value (100:100). Ensure that the same route-target value is configured on all Leaf nodes for MAC VRF to maintain consistency.

```

!
mac vrf VRF1
  rd 1.1.1.1:100
  route-target both 100:100
!
```

Spine Node

1. Configure Interfaces and Loopback

Enable QoS and assign specific IP addresses to loopback interfaces. Configure IP addresses for interfaces connected to each Leaf node.

```

!
qos enable
!
interface ce1/2
  ip address 40.40.40.2/24
!
interface ce1/4
```

```

    ip address 10.10.10.2/24
    !
interface ce24/1
    ip address 30.30.30.2/24
    !
interface ce27/1
    ip address 20.20.20.2/24
    !
interface lo
    ip address 5.5.5.5/32 secondary
    !

```

2. Configure OSPF

In OSPF router mode, set the router ID (5.5.5.5), to match the loopback IP address. Add the loopback network (5.5.5.5/32) and networks (10.10.10.0/24, 20.20.20.0/24, 30.30.30.0/24, and 40.40.40.0/24) connected to Leaf nodes in OSPF area 0. Enable BFD on all OSPF interfaces for faster convergence.

```

!
router ospf 100
    ospf router-id 5.5.5.5
    bfd all-interfaces
    network 5.5.5.5/32 area 0.0.0.0
    network 10.10.10.0/24 area 0.0.0.0
    network 20.20.20.0/24 area 0.0.0.0
    network 30.30.30.0/24 area 0.0.0.0
    network 40.40.40.0/24 area 0.0.0.0
!

```

Configure Switch

Set up an IEEE VLAN bridge, enabling VLANs and associating them with bridge 1. Configure interfaces (xe57, po1, xe46, xe47) to be part of bridge 1, setting them as hybrid ports with VLAN (1000) allowed and egress-tagged enabled. Designate interfaces connected to Leaf nodes (xe46 and xe47) as member ports of po1.

```

!
bridge 1 protocol ieee vlan-bridge
!
vlan database
    vlan-reservation 4000-4094
    vlan 1000 bridge 1 state enable
!
interface po1
    switchport
    bridge-group 1
    switchport mode hybrid
    switchport mode hybrid acceptable-frame-type all
    switchport hybrid allowed vlan add 1000 egress-tagged enable
!
interface xe46
    channel-group 1 mode active
!
interface xe47
    channel-group 1 mode active
!
interface xe57
    switchport
    bridge-group 1

```

```
switchport mode hybrid
switchport mode hybrid acceptable-frame-type all
switchport hybrid allowed vlan add 1000 egress-tagged enable
!
```

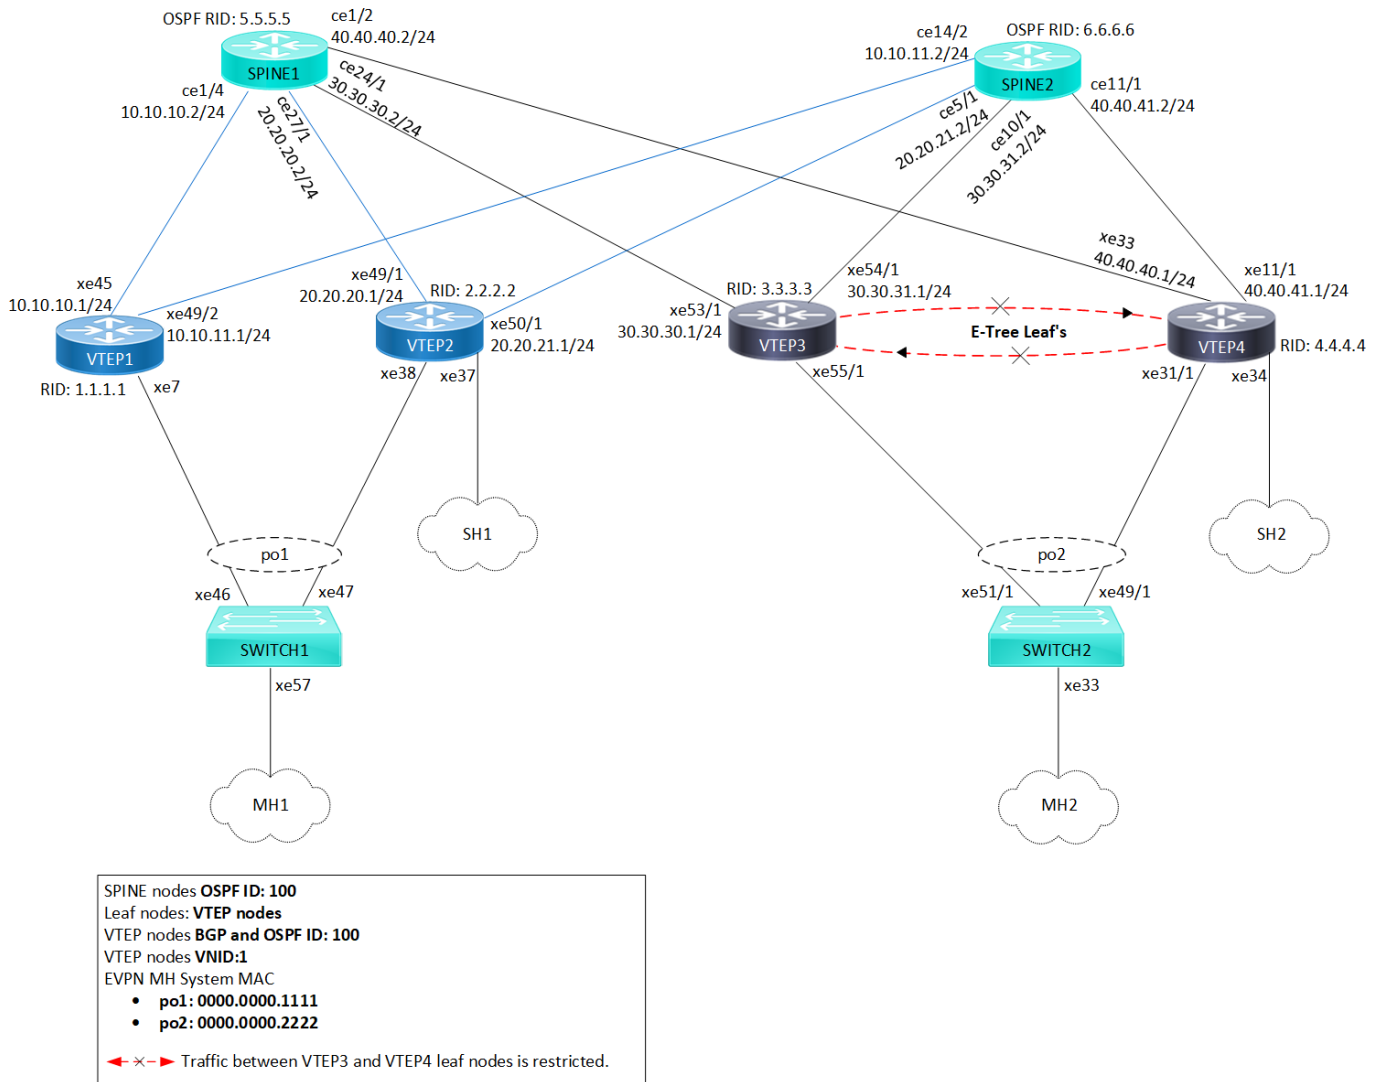
Configuration

Configure various nodes within the topology to set up a VXLAN EVPN E-Tree network.

Topology

The sample topology includes Leaf Nodes (VTEP1, VTEP2, VTEP3, and VTEP4), Spine Nodes (SPINE1 and SPINE2), and Switches (SWITCH1 and SWITCH2).

Within this setup, VTEP1 and VTEP2 are part of Multi-homed group1, where VTEP1 is Multi-homed with po1 (MH1), and VTEP2 is Multi-homed with po1, also featuring an interface as a single home access-if port (SH1). Similarly, VTEP3 and VTEP4 are part of a Multi-homed group2, with VTEP3 Multi-homed with po2 (MH2) and VTEP4 Multi-homed with po2, along with an interface as a single home access-if port (SH2). These nodes connect to Spine nodes (SPINE1 and SPINE2), where all VTEPs are linked. Additionally, SWITCH1 and SWITCH2 are configured to be multihomed to two Leaf nodes each, with SWITCH1 connecting to VTEP1 and VTEP2 and SWITCH2 connecting to VTEP3 and VTEP4.



VXLAN EVPN E-Tree Topology

Note: Before configuring E-Tree, meet all [Prerequisites](#) for the following nodes:

- Leaf nodes: VTEP1, VTEP2, VTEP3, and VTEP4
- Spine nodes: SPINE1 and SPINE2
- Switches: SWITCH1 and SWITCH2

Enable EVPN E-Tree

The following E-Tree configurations applies to the VTEP nodes within the VXLAN network.

1. Enable EVPN E-Tree on VTEP3 and VTEP4 nodes, allowing them to participate in E-Tree functionality within the VXLAN network, controlling traffic and establishing hierarchical connections between Leaf nodes in the network architecture.
 (config)#evpn etree enable
2. Set the ESI hold time (90 seconds) on all VTEP nodes to allow the tunnel to establish during VXLAN initialization before bringing up the ESI. Configure the source VTEP IP address (3.3.3.3) which serves as the global identifier for VXLAN encapsulation and decapsulation within the network, facilitating proper communication and tunnel establishment.

```
(config)#evpn esi hold-time 90
(config)#nvo vxlan vtep-ip-global 3.3.3.3
```

- Define VXLAN identifier (10) with ingress replication and disabled inner VLAN ID (VID) for **E-Tree leaf nodes** (VTEP3 and VTEP4) to support hierarchical connectivity and traffic control within the VXLAN network. This configuration allows for efficient replication of traffic at the ingress point and ensures that inner VLAN IDs are disabled, optimizing the functionality of E-Tree leaf nodes within the network architecture. On the VXLAN tenant node, assign VRF (VRF1) to EVPN-BGP for carrying EVPN routes within the VXLAN network.

```
(config)#nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1
(config-nvo)#exit
```

- Enable port-VLAN mapping (po2) with VLAN ID (1000) to facilitate multi-homed access on all VTEP nodes. Map VXLAN identifier (10) to the access port for VXLAN connectivity.

```
(config)#nvo vxlan access-if port-vlan po2 1000
(config-nvo-acc-if)#map vnid 10
(config-nvo-acc-if)#exit
(config)#commit
```

Validation

Use the show commands described in this section to verify the network for proper VXLAN EVPN E-Tree configuration.

Verify OSPF sessions between the VTEP nodes and the SPINES within the VXLAN network using the `show ip ospf neighbor` command. This command displays OSPF neighbor details, including the state of the OSPF neighbor relationship. A State of Full/DR indicates a fully adjacent and operational state between the routers, confirming proper OSPF connectivity within the network.

```
VTEP1#show ip ospf neighbor
```

```
Total number of full neighbors: 2
```

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

Neighbor ID	Pri	State	Dead Time	Address	Interface	Instance ID
5.5.5.5	1	Full/DR	00:00:32	10.10.10.2	xe45	0
6.6.6.6	1	Full/DR	00:00:30	10.10.11.2	xe49/2	0

Verify the BGP session status between VTEPs, using the `show bgp l2vpn evpn summary` command output. The Up/Down field indicates the duration for which the BGP session has been up or down.

```
VTEP1#show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 9
1 BGP AS-PATH entries
0 BGP community entries
```

Neighbor	V	AS	MsgRcv	MsgSen	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	AD	MACIP	MCAST	ESI	PREFIX-ROUTE
2.2.2.2	4	100	34	28	7	0	0	00:07:37	9	3	4	1	1	0
3.3.3.3	4	100	30	33	8	0	0	00:07:34	6	3	2	1	0	0
4.4.4.4	4	100	31	28	7	0	0	00:07:37	8	3	4	1	0	0

```
Total number of neighbors 3
```

```
Total number of Established sessions 3
```

To validate the BGP L2VPN output on VTEPs and check MAC-IP routes and ESI information, use the `show bgp l2vpn evpn` command output. This command verifies routes with status code `i` (internal) and EVPN route types 2 and 4, displaying detailed information for each VTEP nodes.

```
VTEP1#show bgp l2vpn evpn
BGP table version is 9, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,
l - labeled, S Stale
```


Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevant route information]

1 - Ethernet Auto-discovery Route

2 - MAC/IP Route

3 - Inclusive Multicast Route

4 - Ethernet Segment Route

5 - Prefix Route

Network	Next Hop	Metric	LocPrf	Weight	Path	Peer	Encap
RD[1.1.1.1:100] VRF[VRF1]:							
*> [1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]							
	1.1.1.1	0	100	32768	i	-----	VXLAN
* i	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
*> [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]							
	1.1.1.1	0	100	32768	i	-----	VXLAN
* i	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
* i	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[10]:[10]							
	3.3.3.3	0	100	0	i	3.3.3.3	VXLAN
* i	4.4.4.4	0	100	0	i	4.4.4.4	VXLAN
* i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]							
	3.3.3.3	0	100	0	i	3.3.3.3	VXLAN
* i	3.3.3.3	0	100	0	i	3.3.3.3	VXLAN
* i	4.4.4.4	0	100	0	i	4.4.4.4	VXLAN
* i	4.4.4.4	0	100	0	i	4.4.4.4	VXLAN
*> [2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1000]:[32,100.100.100.1]:[10]							
	1.1.1.1	0	100	32768	i	-----	VXLAN
* i	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
*> [2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1001]:[128,1000::1][10]							
	1.1.1.1	0	100	32768	i	-----	VXLAN
* i	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
* i[2]:[0]:[10]:[48,0000:2000:2000]:[32,200.200.1]:[10]							
	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
* i[2]:[0]:[10]:[48,0000:2000:2001]:[128,2000::1][10]		0	100	0	i	2.2.2.2	VXLAN
* i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3000]:[32,103.103.103.1]:[10]		0	100	0	i	3.3.3.3	VXLAN
* i	3.3.3.3	0	100	0	i	3.3.3.3	VXLAN
* i	4.4.4.4	0	100	0	i	4.4.4.4	VXLAN
* i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3001]:[128,1003::1][10]							
	3.3.3.3	0	100	0	i	3.3.3.3	VXLAN
* i	4.4.4.4	0	100	0	i	4.4.4.4	VXLAN
* i[2]:[0]:[10]:[48,0000:4000:4000]:[32,104.104.104.1]:[10]							
	4.4.4.4	0	100	0	i	4.4.4.4	VXLAN
* i[2]:[0]:[10]:[48,0000:4000:4001]:[128,1004::1][10]		0	100	0	i	4.4.4.4	VXLAN
	4.4.4.4	0	100	0	i	4.4.4.4	VXLAN
*> [3]:[10]:[32,1.1.1.1]							
	1.1.1.1	0	100	32768	i	-----	VXLAN
* i[3]:[10]:[32,2.2.2.2]		0	100	0	i	2.2.2.2	VXLAN
	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
* i[3]:[10]:[32,3.3.3.3]		0	100	0	i	3.3.3.3	VXLAN
	3.3.3.3	0	100	0	i	3.3.3.3	VXLAN
* i[3]:[10]:[32,4.4.4.4]		0	100	0	i	4.4.4.4	VXLAN
	4.4.4.4	0	100	0	i	4.4.4.4	VXLAN
RD[1.1.1.1:64512] VRF[evpn-gvrf-1]:							
*> [1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]							
	1.1.1.1	0	100	32768	i	-----	VXLAN
*> [4]:[00:00:00:00:00:11:11:00:00:00]:[32,1.1.1.1]							
	1.1.1.1	0	100	32768	i	-----	VXLAN
* i[4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]		0	100	0	i	2.2.2.2	VXLAN
	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
RD[2.2.2.2:100]							
*>i[1]:[00:00:00:00:00:11:11:00:00:00]:[10]:[10]							
	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
*>i[1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]							
	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
*>i[2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1000]:[32,100.100.100.1]:[10]		0	100	0	i	2.2.2.2	VXLAN
	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN
*>i[2]:[00:00:00:00:00:11:11:00:00:00]:[10]:[48,0000:1000:1001]:[128,1000::1][10]		0	100	0	i	2.2.2.2	VXLAN
	2.2.2.2	0	100	0	i	2.2.2.2	VXLAN

```

*>i[2]:[0]:[10]:[48,0000:2000:2000]:[32,200.200.200.1]:[10]
    2.2.2.2          0          100      0    i  2.2.2.2      VXLAN
*>i[2]:[0]:[10]:[48,0000:2000:2001]:[128,2000::1][10]
    2.2.2.2          0          100      0    i  2.2.2.2      VXLAN
*>i[3]:[10]:[32,2.2.2.2]
    2.2.2.2          0          100      0    i  2.2.2.2      VXLAN

RD[2.2.2.2:64512]
*>i[1]:[00:00:00:00:00:11:11:00:00:00]:[4294967295]:[0]
    2.2.2.2          0          100      0    i  2.2.2.2      VXLAN
*>i[4]:[00:00:00:00:00:11:11:00:00:00]:[32,2.2.2.2]
    2.2.2.2          0          100      0    i  2.2.2.2      VXLAN

RD[3.3.3.3:100]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[10]:[10]
    3.3.3.3          0          100      0    i  3.3.3.3      VXLAN
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
    3.3.3.3          0          100      0    i  3.3.3.3      VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3000]:[32,103.103.103.1]:[10]
    3.3.3.3          0          100      0    i  3.3.3.3      VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3001]:[128,1003::1][10]
    3.3.3.3          0          100      0    i  3.3.3.3      VXLAN
*>i[3]:[10]:[32,3.3.3.3]
    3.3.3.3          0          100      0    i  3.3.3.3      VXLAN

RD[3.3.3.3:64512]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
    3.3.3.3          0          100      0    i  3.3.3.3      VXLAN

RD[4.4.4.4:100]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[10]:[10]
    4.4.4.4          0          100      0    i  4.4.4.4      VXLAN
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
    4.4.4.4          0          100      0    i  4.4.4.4      VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3000]:[32,103.103.103.1]:[10]
    4.4.4.4          0          100      0    i  4.4.4.4      VXLAN
*>i[2]:[00:00:00:00:00:22:22:00:00:00]:[10]:[48,0000:3000:3001]:[128,1003::1][10]
    4.4.4.4          0          100      0    i  4.4.4.4      VXLAN
*>i[2]:[0]:[10]:[48,0000:4000:4000]:[32,104.104.104.1]:[10]
    4.4.4.4          0          100      0    i  4.4.4.4      VXLAN
*>i[2]:[0]:[10]:[48,0000:4000:4001]:[128,1004::1][10]
    4.4.4.4          0          100      0    i  4.4.4.4      VXLAN
*>i[3]:[10]:[32,4.4.4.4]
    4.4.4.4          0          100      0    i  4.4.4.4      VXLAN

RD[4.4.4.4:64512]
*>i[1]:[00:00:00:00:00:22:22:00:00:00]:[4294967295]:[0]
    4.4.4.4          0          100      0    i  4.4.4.4      VXLAN

```

Total number of prefixes 42

Validate the LAG interfaces (po1 and po2) are up for MH1 and MH2 by reviewing the show etherchannel summary output. Check the Link and sync fields, where link displays the port channel interface and ID number, and sync indicates whether MAC address synchronization is enabled to forward Layer 3 packets arriving on these interfaces.

```

VTEP1#show etherchannel summary
  Aggregator po1 100001
  Aggregator Type: Layer2
  Admin Key: 0001 - Oper Key 0001
  Link: xe7 (5005) sync: 1

```

Validate the status of NVO VXLAN on VTEPs by examining the output of the show nvo vxlan command. The DF-Status field displays the forwarding status of VXLAN tunnels as a Designated Forwarder (DF) or Non-Designated Forwarder (Non-DF).

```

VTEP1#show nvo vxlan
VXLAN Information
=====

```

Codes: NW - Network Port
 AC - Access Port
 (u) - Untagged

VNID	VNI-Name	VNI-Type	Type	Interface	ESI	VLAN	DF-Status	Src-Addr	Dst-Addr
10	----	L2	NW	----	----	----	----	1.1.1.1	4.4.4.4
10	----	L2	NW	----	----	----	----	1.1.1.1	3.3.3.3
10	----	L2	NW	----	----	----	----	1.1.1.1	2.2.2.2
10	----	--	AC	po1	00:00:00:00:00:11:11:00:00:00	1000	DF	----	----

Total number of entries are 4

VTEP2#show nvo vxlan
 VXLAN Information

Codes: NW - Network Port
 AC - Access Port
 (u) - Untagged

VNID	VNI-Name	VNI-Type	Type	Interface	ESI	VLAN	DF-Status	Src-Addr	Dst-Addr
10	----	L2	NW	----	----	----	----	2.2.2.2	4.4.4.4
10	----	L2	NW	----	----	----	----	2.2.2.2	1.1.1.1
10	----	L2	NW	----	----	----	----	2.2.2.2	3.3.3.3
10	----	--	AC	xe37	--- Single Homed Port ---	1000	----	----	----
10	----	--	AC	po1	00:00:00:00:00:11:11:00:00:00	1000	NON-DF	----	----

Total number of entries are 5

VTEP3#show nvo vxlan
 VXLAN Information

Codes: NW - Network Port
 AC - Access Port
 (u) - Untagged

VNID	VNI-Name	VNI-Type	Type	Interface	ESI	VLAN	DF-Status	Src-Addr	Dst-Addr
10	----	L2	NW	----	----	----	----	3.3.3.3	2.2.2.2
10	----	L2	NW	----	----	----	----	3.3.3.3	1.1.1.1
10	----	L2	NW	----	----	----	----	3.3.3.3	4.4.4.4
10	----	--	AC	po2	00:00:00:00:00:22:22:00:00:00	1000	DF	----	----

Total number of entries are 4

VTEP4#show nvo vxlan
 VXLAN Information

Codes: NW - Network Port
 AC - Access Port
 (u) - Untagged

VNID	VNI-Name	VNI-Type	Type	Interface	ESI	VLAN	DF-Status	Src-Addr	Dst-Addr
10	----	L2	NW	----	----	----	----	4.4.4.4	2.2.2.2
10	----	L2	NW	----	----	----	----	4.4.4.4	3.3.3.3
10	----	L2	NW	----	----	----	----	4.4.4.4	1.1.1.1
10	----	--	AC	xe34	--- Single Homed Port ---	1000	----	----	----
10	----	--	AC	po2	00:00:00:00:00:22:22:00:00:00	1000	NON-DF	----	----

Total number of entries are 5

Validate the NVO VXLAN tunnel status on VTEPs by reviewing the output of the show nvo vxlan tunnel command. The Status field indicates the current status of each tunnel. In this case, all three tunnels between VTEPs and their respective destinations are marked as Installed, confirming that these tunnels are successfully established and operating.

VTEP1#show nvo vxlan tunnel
 VXLAN Network tunnel Entries

Source	Destination	Status	Up/Down	Update
--------	-------------	--------	---------	--------

```
=====
1.1.1.1      4.4.4.4      Installed      00:02:26      00:01:58
1.1.1.1      3.3.3.3      Installed      00:02:26      00:01:55
1.1.1.1      2.2.2.2      Installed      00:02:25      00:01:55
```

Total number of entries are 3

Validate the VXLAN access interface status on VTEPs by examining the output of the `show nvo vxlan access-if brief` command. The `up admin` and `link status` confirms that the access port associated with VXLAN is active and functioning properly on the VTEP nodes.

```
VTEP1#show nvo vxlan access-if brief
```

```

          Inner
Interface  Vlan      vlan  Ifindex  Vnid      Admin  Link
          status status
-----
po1        1000     ---   0x7a120   10        up     up
```

Total number of entries are 1

Static MAC-IP Advertisement

Configure static MAC-IP advertisement through SH and MH VTEPs from Root and Leaf nodes. Advertise static MAC addresses for IPv4 and IPv6 from MH1, MH2, SH1, and SH2 VTEPs. Ensure that VTEP1 and VTEP2 in MH1 have the same MAC addresses configured under the port-channel access port. Symmetrical configurations between MH VTEPs should be maintained.

Configure MH1 and MH2 VTEPs

Configure static MAC addresses for IPv4 (100.100.100.1) and IPv6 (1000::1) under the VXLAN MH access-port (po1) with VLAN ID (1000). Ensure that identical MAC addresses are set up within the MH1-VTEPs for advertisement. Apply similar configurations to MH2-VTEPs for static MAC-IP advertisement.

```
!
nvo vxlan access-if port-vlan po1 1000
  map vnid 10
  mac 0000.1000.1000 ip 100.100.100.1
  mac 0000.1000.1001 ipv6 1000::1
!
```

Configure SH1 and SH2 VTEPs

Configure static MAC addresses for IPv4 (200.200.200.1) and IPv6 (2000::1) under the VXLAN SH access-port (xe37) with VLAN ID (1000) on SH1 (VTEP2). This setup ensures that SH1 advertises these static MAC addresses over the specified VXLAN access-port. Repeat similar configurations for SH2 (VTEP4) using different static MAC addresses for both IPv4 and IPv6.

```
!
nvo vxlan access-if port-vlan xe37 1000
  map vnid 10
  mac 0000.2000.2000 ip 200.200.200.1
  mac 0000.2000.2001 ipv6 2000::1
!
```

Validation

Verify the MAC table entries on MH VTEPs (MH1 and MH2) and the SH VTEPs (VTEP2 and VTEP4). The MAC addresses are advertised using the ESI values from VTEP1 and VTEP2 for MH1, and from VTEP3 and VTEP4 for MH2. Additionally, verify the VTEP IP addresses associated with SH VTEP2 and VTEP4 for MAC advertisement.

In the output of the `show nvo vxlan mac-table` command on all VTEP nodes, the MAC entries advertised from Leaf VTEPs will have the `LeafFlag` field status `set`.

Note:

- MAC IPv4 or IPv6 configured under SH Leaf VTEP access port will be advertised to the Root VTEP and other Leaf VTEPs.
- MAC IPv4 or IPv6 configured under an MH Leaf VTEP access port must be symmetric and will be advertised to both the Root VTEP and other leaf VTEPs.
- MAC IPv4 or IPv6 configured under either SH or MH Root VTEP will be advertised to both the Root VTEP and the Leaf VTEPs.
- The Leaf-to-Leaf communication will display MAC status and tunnel status per VNI as Leaf type. The MAC will be in the discard state in the BCM shell.

VTEP1#show nvo vxlan mac-table

```
=====
                                VXLAN MAC Entries
=====
VNID Interface VlanId  In-VlanId Mac-Addr          VTEP-Ip/ESI          Type  Status  MAC move AccessPortDesc LeafFlag
-----
10  po1      1000    ----    0000.1000.1000    00:00:00:00:00:11:11:00:00:00 Static Local  ----- 0 -----  ----
10  po1      1000    ----    0000.1000.1001    00:00:00:00:00:11:11:00:00:00 Static Local  ----- 0 -----  ----
10  ----     ----     ----    0000.2000.2000    2.2.2.2              Static Remote ----- 0 -----  ----
10  ----     ----     ----    0000.2000.2001    2.2.2.2              Static Remote ----- 0 -----  ----
10  ----     ----     ----    0000.3000.3000    00:00:00:00:00:22:22:00:00:00 Static Remote ----- 0 -----  set
10  ----     ----     ----    0000.3000.3001    00:00:00:00:00:22:22:00:00:00 Static Remote ----- 0 -----  set
10  ----     ----     ----    0000.4000.4000    4.4.4.4              Static Remote ----- 0 -----  set
10  ----     ----     ----    0000.4000.4001    4.4.4.4              Static Remote ----- 0 -----  set
```

Total number of entries are : 8

VTEP3#show nvo vxlan mac-table

```
=====
                                VXLAN MAC Entries
=====
VNID Interface VlanId  In-VlanId Mac-Addr          VTEP-Ip/ESI          Type  Status  MAC move AccessPortDesc LeafFlag
-----
10  ----     ----     ----    0000.1000.1000    00:00:00:00:00:11:11:00:00:00 Static Remote ----- 0 -----  ----
10  ----     ----     ----    0000.1000.1001    00:00:00:00:00:11:11:00:00:00 Static Remote ----- 0 -----  ----
10  ----     ----     ----    0000.2000.2000    2.2.2.2              Static Remote ----- 0 -----  ----
10  ----     ----     ----    0000.2000.2001    2.2.2.2              Static Remote ----- 0 -----  ----
10  po2      1000    ----    0000.3000.3000    00:00:00:00:00:22:22:00:00:00 Static Local  ----- 0 -----  set
10  po2      1000    ----    0000.3000.3001    00:00:00:00:00:22:22:00:00:00 Static Local  ----- 0 -----  set
10  ----     ----     ----    0000.4000.4000    4.4.4.4              Static Remote ----- 0 -----  set
10  ----     ----     ----    0000.4000.4001    4.4.4.4              Static Remote ----- 0 -----  set
```

Total number of entries are : 8

Use the `show nvo vxlan arp-cache` command to verify the Address Resolution Protocol (ARP) cache information on all VTEP nodes. This command displays entries that map IPv4 addresses to MAC addresses within the specified VXLAN VNID network.

VTEP1#show nvo vxlan arp-cache

VXLAN ARP-CACHE Information

```
=====
VNID      Ip-Addr          Mac-Addr          Type          Age-Out  Retries-Left
-----
10        100.100.100.1    0000.1000.1000    Static        Local    -----
```

```

10      103.103.103.1    0000.3000.3000 Static      Remote      ----
10      104.104.104.1    0000.4000.4000 Static      Remote      ----
10      200.200.200.1    0000.2000.2000 Static      Remote      ----
Total number of entries are 4

```

```

VTEP3#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information

```

```

=====
VNID      Ip-Addr          Mac-Addr          Type          Age-Out      Retries-Left
-----
10      100.100.100.1    0000.1000.1000 Static Remote      ----
10      103.103.103.1    0000.3000.3000 Static Local      ----
10      104.104.104.1    0000.4000.4000 Static Remote      ----
10      200.200.200.1    0000.2000.2000 Static Remote      ----
Total number of entries are 4

```

Use the `show nvo vxlan nd-cache` command to verify the Neighbor Discovery (ND) cache information on all VTEP nodes. This command displays entries that map IPv6 addresses to MAC addresses within the specified VXLAN VNID network.

```

VTEP1#show nvo vxlan nd-cache
VXLAN ND-CACHE Information

```

```

=====
VNID      Ip-Addr          Mac-Addr          Type          Age-Out      Retries-Left
-----
10      1000:::1         0000.1000.1001 Static Local      ----
10      1003:::1         0000.3000.3001 Static Remote      ----
10      1004:::1         0000.4000.4001 Static Remote      ----
10      2000:::1         0000.2000.2001 Static Remote      ----
Total number of entries are 4

```

```

VTEP3#show nvo vxlan nd-cache
VXLAN ND-CACHE Information

```

```

=====
VNID      Ip-Addr          Mac-Addr          Type          Age-Out      Retries-Left
-----
10      1000:::1         0000.1000.1001 Static Remote      ----
10      1003:::1         0000.3000.3001 Static Local      ----
10      1004:::1         0000.4000.4001 Static Remote      ----
10      2000:::1         0000.2000.2001 Static Remote      ----
Total number of entries are 4

```

Network Topology Snippet Configurations

Here are the snippet configurations for all nodes in the given network topology.

VTEP1

```

!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!

```

```
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
mac vrf VRF1
  rd 1.1.1.1:100
  route-target both 100:100
!
nvo vxlan vtep-ip-global 1.1.1.1
!
nvo vxlan id 10 ingress-replication inner-vid-disabled
  vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
interface po1
  switchport
  evpn multi-homed system-mac 0000.0000.1111
!
interface lo
  ip address 1.1.1.1/32 secondary
!
interface xe7
  channel-group 1 mode active
!
interface xe45
  ip address 10.10.10.1/24
!
interface xe49/2
  ip address 10.10.11.1/24
!
exit
!

router ospf 100
  ospf router-id 1.1.1.1
  bfd all-interfaces
  network 1.1.1.1/32 area 0.0.0.0
  network 10.10.10.0/24 area 0.0.0.0
  network 10.10.11.0/24 area 0.0.0.0
!
router bgp 100
  bgp router-id 1.1.1.1
  neighbor 2.2.2.2 remote-as 100
  neighbor 3.3.3.3 remote-as 100
  neighbor 4.4.4.4 remote-as 100
  neighbor 2.2.2.2 update-source lo
  neighbor 2.2.2.2 advertisement-interval 0
  neighbor 3.3.3.3 update-source lo
  neighbor 3.3.3.3 advertisement-interval 0
  neighbor 4.4.4.4 update-source lo
  neighbor 4.4.4.4 advertisement-interval 0
```

```

!
address-family l2vpn evpn
neighbor 2.2.2.2 activate
neighbor 3.3.3.3 activate
neighbor 4.4.4.4 activate
exit-address-family
!
exit
!
nvo vxlan access-if port-vlan pol 1000
map vnid 10
mac 0000.1000.1000 ip 100.100.100.1
mac 0000.1000.1001 ipv6 1000::1
!

```

VTEP2

```

!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
mac vrf VRF1
rd 2.2.2.2:100
route-target both 100:100
!
nvo vxlan vtep-ip-global 2.2.2.2
!
nvo vxlan id 10 ingress-replication inner-vid-disabled
vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
interface pol
switchport
evpn multi-homed system-mac 0000.0000.1111
!
interface lo
ip address 2.2.2.2/32 secondary
!
interface xe38
channel-group 1 mode active
!
interface xe49/1
ip address 20.20.20.1/24
!
interface xe50/1
ip address 20.20.21.1/24

```



```

!
  exit
!

router ospf 100
  ospf router-id 2.2.2.2
  bfd all-interfaces
  network 2.2.2.2/32 area 0.0.0.0
  network 20.20.20.0/24 area 0.0.0.0
  network 20.20.21.0/24 area 0.0.0.0
!
router bgp 100
  bgp router-id 2.2.2.2
  neighbor 1.1.1.1 remote-as 100
  neighbor 3.3.3.3 remote-as 100
  neighbor 4.4.4.4 remote-as 100
  neighbor 1.1.1.1 update-source lo
  neighbor 1.1.1.1 advertisement-interval 0
  neighbor 3.3.3.3 update-source lo
  neighbor 3.3.3.3 advertisement-interval 0
  neighbor 4.4.4.4 update-source lo
  neighbor 4.4.4.4 advertisement-interval 0
!
  address-family l2vpn evpn
  neighbor 1.1.1.1 activate
  neighbor 3.3.3.3 activate
  neighbor 4.4.4.4 activate
  exit-address-family
!
  exit
!
nvo vxlan access-if port-vlan xe37 1000
  map vnid 10
  mac 0000.2000.2000 ip 200.200.200.1
  mac 0000.2000.2001 ipv6 2000::1
!
nvo vxlan access-if port-vlan po1 1000
  map vnid 10
  mac 0000.1000.1000 ip 100.100.100.1
  mac 0000.1000.1001 ipv6 1000::1
!

```

VTEP3

```

!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
evpn etree enable
!

```

```
mac vrf VRF1
  rd 3.3.3.3:100
  route-target both 100:100
!
nvo vxlan vtep-ip-global 3.3.3.3
!
nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
  vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
interface po2
  switchport
  evpn multi-homed system-mac 0000.0000.2222
!
interface lo
  ip address 3.3.3.3/32 secondary
!
interface xe53/1
  ip address 30.30.30.1/24
!
interface xe54/1
  ip address 30.30.31.1/24
!
interface xe55/1
  channel-group 2 mode active
!
exit
!
router ospf 100
  ospf router-id 3.3.3.3
  bfd all-interfaces
  network 3.3.3.3/32 area 0.0.0.0
  network 30.30.30.0/24 area 0.0.0.0
  network 30.30.31.0/24 area 0.0.0.0
!
router bgp 100
  bgp router-id 3.3.3.3
  neighbor 1.1.1.1 remote-as 100
  neighbor 2.2.2.2 remote-as 100
  neighbor 4.4.4.4 remote-as 100
  neighbor 1.1.1.1 update-source lo
  neighbor 1.1.1.1 advertisement-interval 0
  neighbor 2.2.2.2 update-source lo
  neighbor 2.2.2.2 advertisement-interval 0
  neighbor 4.4.4.4 update-source lo
  neighbor 4.4.4.4 advertisement-interval 0
!
  address-family l2vpn evpn
  neighbor 1.1.1.1 activate
  neighbor 2.2.2.2 activate
  neighbor 4.4.4.4 activate
  exit-address-family
!
exit
!
!
```

```
nvo vxlan access-if port-vlan po2 1000
  map vnid 10
  mac 0000.3000.3000 ip 103.103.103.1
  mac 0000.3000.3001 ipv6 1003::1
!
```

VTEP4

```
!
hardware-profile filter vxlan enable
hardware-profile filter vxlan-mh enable
!
nvo vxlan enable
!
evpn esi hold-time 90
!
evpn vxlan multihoming enable
!
evpn etree enable
!
mac vrf VRF1
  rd 4.4.4.4:100
  route-target both 100:100
!
nvo vxlan vtep-ip-global 4.4.4.4
!
nvo vxlan id 10 ingress-replication inner-vid-disabled etree-leaf
  vxlan host-reachability-protocol evpn-bgp VRF1
!
qos enable
!
interface po2
  switchport
  evpn multi-homed system-mac 0000.0000.2222
!
interface lo
  ip address 4.4.4.4/32 secondary
!
interface xe11/1
  ip address 40.40.41.1/24
!
interface xe31/1
  channel-group 2 mode active
!
interface xe33
  ip address 40.40.40.1/24
!
interface xe34
  switchport
!
exit
!
router ospf 100
  ospf router-id 4.4.4.4
  bfd all-interfaces
  network 4.4.4.4/32 area 0.0.0.0
```

```
network 40.40.40.0/24 area 0.0.0.0
network 40.40.41.0/24 area 0.0.0.0
!
router bgp 100
  bgp router-id 4.4.4.4
  neighbor 1.1.1.1 remote-as 100
  neighbor 2.2.2.2 remote-as 100
  neighbor 3.3.3.3 remote-as 100
  neighbor 1.1.1.1 update-source lo
  neighbor 1.1.1.1 advertisement-interval 0
  neighbor 2.2.2.2 update-source lo
  neighbor 2.2.2.2 advertisement-interval 0
  neighbor 3.3.3.3 update-source lo
  neighbor 3.3.3.3 advertisement-interval 0
!
  address-family l2vpn evpn
  neighbor 1.1.1.1 activate
  neighbor 2.2.2.2 activate
  neighbor 3.3.3.3 activate
  exit-address-family
!
exit
!
nvo vxlan access-if port-vlan xe34 1000
  map vnid 10
  mac 0000.4000.4000 ip 104.104.104.1
  mac 0000.4000.4001 ipv6 1004::1
!
nvo vxlan access-if port-vlan po2 1000
  map vnid 10
  mac 0000.3000.3000 ip 103.103.103.1
  mac 0000.3000.3001 ipv6 1003::1
!
```

SPINE1

```
!
qos enable
!
interface ce1/2
  ip address 40.40.40.2/24
!
interface ce1/4
  ip address 10.10.10.2/24
!
interface ce24/1
  ip address 30.30.30.2/24
!
interface ce27/1
  ip address 20.20.20.2/24
!
interface lo
  ip address 5.5.5.5/32 secondary
!
exit
!
```

```
router ospf 100
  ospf router-id 5.5.5.5
  bfd all-interfaces
  network 5.5.5.5/32 area 0.0.0.0
  network 10.10.10.0/24 area 0.0.0.0
  network 20.20.20.0/24 area 0.0.0.0
  network 30.30.30.0/24 area 0.0.0.0
  network 40.40.40.0/24 area 0.0.0.0
!
```

SPINE2

```
!
qos enable
!
interface ce5/1
  ip address 20.20.21.2/24
!
interface ce10/1
  ip address 30.30.31.2/24
!
interface ce11/1
  ip address 40.40.41.2/24
!
interface ce14/2
  ip address 10.10.11.2/24
!
interface lo
  ip address 6.6.6.6/32 secondary
!
exit
!
router ospf 100
  ospf router-id 6.6.6.6
  bfd all-interfaces
  network 6.6.6.6/32 area 0.0.0.0
  network 10.10.11.0/24 area 0.0.0.0
  network 20.20.21.0/24 area 0.0.0.0
  network 30.30.31.0/24 area 0.0.0.0
  network 40.40.41.0/24 area 0.0.0.0
!
```

SWITCH1

```
!
bridge 1 protocol ieee vlan-bridge
!
vlan database
  vlan-reservation 4000-4094
  vlan 1000 bridge 1 state enable
!
interface po1
  switchport
  bridge-group 1
  switchport mode hybrid
  switchport mode hybrid acceptable-frame-type all
```

```

    switchport hybrid allowed vlan add 1000 egress-tagged enable
    !
interface xe46
    channel-group 1 mode active
    !
interface xe47
    channel-group 1 mode active
    !
interface xe57
    switchport
    bridge-group 1
    switchport mode hybrid
    switchport mode hybrid acceptable-frame-type all
    switchport hybrid allowed vlan add 1000 egress-tagged enable
    !
    exit
    !

```

SWITCH2

```

    !
bridge 1 protocol ieee vlan-bridge
    !
vlan database
    vlan-reservation 4000-4094
    vlan 1000 bridge 1 state enable
    !
interface po2
    switchport
    bridge-group 1
    switchport mode hybrid
    switchport mode hybrid acceptable-frame-type all
    switchport hybrid allowed vlan add 1000 egress-tagged enable
    !
interface xe33
    switchport
    bridge-group 1
    switchport mode hybrid
    switchport mode hybrid acceptable-frame-type all
    switchport hybrid allowed vlan add 1000 egress-tagged enable
    !
interface xe49/1
    channel-group 2 mode active
    !
interface xe51/1
    channel-group 2 mode active
    !
    exit
    !

```

Implementation Examples

Here is an example scenario and a solution for implementing EVPN E-Tree.

Scenario 1: Specific traffic isolation and control measures are essential in a network of EVPN L2VPN services or instances. Within a broadcast domain, services communicating with each other may result in flooding BUM traffic to all services within the domain. Moreover, hosts are learned and advertised between different sites/services.

Use Case 1: Implementing an EVPN E-Tree solution defines the network topology with distinct Root and Leaf classifications, BUM traffic flooding can be minimized, and traffic isolation can be achieved. This ensures efficient communication between services while preventing unnecessary traffic propagation and maintaining network integrity.

Scenario 2: An Internet Service Provider (ISP) provides services to multiple subscribers and aims to facilitate communication with them. However, the ISP needs to ensure that subscribers exclusively communicate with the ISP and not among themselves.

Use Case 2: Implementing EVPN E-Tree is essential to fulfill this requirement. By categorizing ISP services as Root and subscribers as Leaf, traffic isolation can be enforced. This configuration enables the ISP to communicate with subscribers while preventing inter-subscriber communication. As a result, network security is enhanced, and the ISP maintains control over communication within its network.

E-Tree CLI Commands

The EVPN E-Tree introduces the following configuration commands in OcNOS.

evpn etree

Use this command to enable E-Tree functionality within the EVPN configuration.

Command Syntax

```
evpn etree enable
```

Parameters

None

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.1.

Example

The following example illustrates how to activate E-Tree functionality for EVPN:

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable
```

Revised CLI Commands

The following is the revised command for configuring VXLAN EVPN E-Tree

nvo vxlan id

- The existing syntax now includes the newly added parameter for E-Tree, namely `etree-leaf`.
- The command `nvo vxlan id <VNID> ingress-replication inner-vid-disabled etree-leaf` allows users to tailor VXLAN behavior on a network device, specifying VXLAN parameters and indicating its participation as a leaf node in an E-Tree deployment. For more details, refer to the [nvo vxlan id](#) command in the [VXLAN Commands](#) chapter in the *OcNOS VXLAN Guide*.

Troubleshooting

1. When traffic, whether unicast (UC) or broadcast, is passed to the Intra Leaf site:
 - Check the sub-interface or physical interface counters to monitor traffic throughput and potential issues.
 - Verify the Leaf status of the corresponding VNI to ensure proper functionality.
 - Use packet sniffing tools to analyze packets in the egress direction for any anomalies or errors.
 - MAC entries learned via leaf access port should include the `set` keyword in the MAC table output.
2. If UC traffic is routed within inter-PE leaf sites:
 - Check the Leaf status of the VNI at both participating PE devices to confirm operational status.
 - Check if the advertised MAC is in discard or non-discard status using the `show mac table` command and `12 show` in the BCM shell.
3. Investigate UC traffic drops from the Root to MH Leaf PE:
 - Check if MAC addresses are not installed in discard status within the MH peer's access port. This status could indicate issues with MAC learning or forwarding.
4. Evaluate traffic between Root and Leaf:
 - Confirm the establishment of both UC and BUM tunnels.
 - Ensure that unicast MAC addresses are not marked with a discard status in the MAC table.
5. Validate the exchange of routes between two BGP L2VPN peers:
 - Monitor BGP (Border Gateway Protocol) sessions to verify successful route exchange and propagation between the peers.
6. Convergence: Assess convergence by checking BFD configuration between BGP sessions.

Glossary

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

Key Terms/Acronym	Description
Ethernet VPN Ethernet-Tree (EVPN E-Tree)	A networking solution designed to manage communication within broadcast domains, incorporating redundancy through multi-homing in a network. It optimizes traffic routing and control, categorizing network nodes based on predefined definitions of EVPN Instances as Leaf or Root, allowing or restricting communication between them.

Virtual Extensible LAN (VXLAN)	A technology that provides encapsulation techniques to create virtualized Layer 2 networks over Layer 3 infrastructure, facilitating scalable and flexible network designs.
Ethernet Virtual Private Network (EVPN)	A Layer 2 VPN technology that extends Ethernet services across data centers and wide-area networks using BGP.
Multi-homing (MH)	The ability of a device to connect to multiple network segments simultaneously to increase network availability and redundancy.
Provider Edge (PE) Node	A device at the edge of a service provider network that connects to customer premises equipment (CE) and participates in providing services to customers.
Leaf Node	In the context of EVPN E-Tree, a network node categorized to handle communication within specific broadcast domains and may connect to Root nodes.
Root Node	A network node within EVPN E-Tree that serves as the central point of communication and handles BUM traffic distribution.
Ethernet Segment Identifier (ESI)	A unique identifier used to identify Ethernet segments within a VXLAN network.

Improved Routing

This section describes the new feature for improved network routing introduced in the Release 6.5.2.

- [ISIS Multi Topology](#)

CHAPTER 1 ISIS Multi Topology

Overview

Intermediate System to Intermediate System (ISIS) is a link-state routing protocol commonly used in large-scale service provider networks and enterprise networks. By default, ISIS is in a single topology with no separate Shortest Path First (SPF) process to differentiate between IPv4 and IPv6 topologies. If the topology in IPv6 is different from IPv4, the routing calculation encounters a problem as the routes are evaluated and chosen based on the common topology.

Multi Topology (MT) is a mechanism to run a set of independent IP topologies within a single ISIS domain. This means, both IPv4 and IPv6 have different topologies in the network and two SPF processes are run to find the route to each IPv4 and IPv6 destination independently.

Feature Characteristics

The main characteristics of ISIS Multi Topology are as follows:

- Enables ISIS to maintain separate topologies for IPv4 and IPv6 within the same ISIS area or domain.
- Allows routers in the ISIS area (for Level 1 routing) or domain (for Level 2 routing) to support both IPv4 and IPv6 address families.
- Performs multiple SPF calculations for each configured topology.
- Defines new Type-Length-Value (TLV) encodings called Multi Topology TLV (MT TLV). It is used to advertise the multiple topologies supported by the routers and contains information about the topology, including the ID (MTID), flags, and MT metric.
 - MT TLV (229): Capability TLV advertised in Hello packets.
 - MT intermediate system TLV (222): Extended TLV that describes the adjacency between nodes once the adjacency is formed.
 - MT IPV6 reachability TLV (237): Reachability TLV that gives information on IPv6 routing.

Benefits

The key benefits of ISIS Multi Topology are as follows:

- Enables the ability to make changes to the IPv6 topology without affecting the IPv4 topology, and vice-versa.
- Leverages common adjacency and database tables.
- Provides an independent SPF process for IPv4 and IPv6.

Prerequisite

To enable ISIS Multi Topology on OcNOS devices, wide metric configuration is mandatory.

Configuration

To set up Multi Topology in ISIS, the configuration is as shown below:

Topology

This topology diagram consists of five routers (R1, R2, R3, R4 and R5).

It has both ISIS IPv4 and IPv6 routing enabled, except the link between R2 and R4 which has only IPv6 enabled.

In Single Topology, router R1 receives the information and calculates a SPF tree. To reach 5.5.5.5 (R5 IPv4), it takes the path R1-> R2 -> R4 ->R5. However, it fails since R2 to R4 is solely an IPv6 path. Since the same SPF tree is used for both IPv4 and IPv6 in R1, it considers the link between R2 -> R4 as the shortest path instead of R2 -> R3 -> R4.

On enabling Multi Topology on all the routers, SPF trees are calculated separately for IPv4 and IPv6 routing. This means, to reach from R1 to R5, IPv4 takes the R1 -> R2 -> R3 -> R4 -> R5 path and IPv6 takes the R1 -> R2 -> R4 -> R5 path.




Figure 1-10: ISIS Multi Topology

To configure multi topology on the routers R1, R2, R3, R4 and R5, follow the steps mentioned below:

The first set of configuration commands enable single topology per router in preparation for the interface to implement multi topology.

R1

1. Enter configure mode followed by interface mode on loopback interface.

```
#configure terminal
R1(config)#int lo
```

2. Configure the IP address for the interface.

```
R1(config-if)# ip add 1.1.1.1/32 secondary
R1(config-if)# ipv6 address 1111::11/128
```

3. Include the interface in the router's ISIS 1 instance.

```
R1(config-if)# ip router isis 1
R1(config-if)# ipv6 router isis 1
```

4. Exit the interface mode.

```
R1(config-if)# exit
```

5. Enter the interface configuration mode.

```
R1(config)#int xe22
```

6. Configure the IP address for the interface.

```
R1(config -if)# ip address 10.1.1.1/24
R1(config -if)# ipv6 address 1001::1/64
```

7. Include the interface in the router's ISIS 1 instance.

```
R1(config -if)# ip router isis 1
R1(config -if)# ipv6 router isis 1
```

8. Exit the interface mode.

```
R1(config -if)# exit
```

9. Set the routing process ID as 1.

```
R1(config)# router isis 1
```

10. Configure the IS type as level-1.

```
R1(config-router)# is-type level-1
```

11. Configure wide metric-style.

```
R1(config-router)# metric-style wide
```

12. Enable dynamic host name under ISIS process.

```
R1(config-router)# dynamic-hostname
```

13. Enable BFD in all the interfaces.

```
R1(config-router)# bfd all-interfaces
```

14. Configure Network Entity Title (NET).

```
R1(config-router)# net 49.0000.0000.0001.00
```

15. Commit the candidate configuration to the running configuration.

```
R1(config-router)# commit
```

R2

1. Enter configure mode followed by interface mode for loopback interface.

```
#configure terminal
R2(config)#int lo
```

2. Configure the IP address for the interface.

```
R2(config -if)# ip add 2.2.2.2/32 secondary
R2(config -if)# ipv6 address 2222::22/128
```

3. Include the interface in the router's ISIS 1 instance.

```
R2(config -if)# ip router isis 1
R2(config -if)# ipv6 router isis 1
```

4. Exit the interface mode.

```
R2(config -if)# exit
```

5. Enter the interface configuration mode.

```
R2(config)#int xe12
```

6. Configure the IP address for the interface.

```
R2(config -if)# ip address 10.1.1.2/24
R2(config -if)# ipv6 address 1001::2/64
```

7. Include the interface in the router's ISIS 1 instance.

```
R2(config -if)# ip router isis 1
R2(config -if)# ipv6 router isis 1
```

8. Exit the interface mode.

```
R2(config -if)# exit
```

9. Enter the interface configuration mode

```
R2(config)#int xe24
```

10. Configure the IP address for the interface.

```
R2(config -if)# ip address 20.1.1.1/24
R2(config -if)# ipv6 address 2001::1/64
```

11. Include the interface in the router's ISIS 1 instance.

```
R2(config -if)# ipv6 router isis 1
```

12. Exit the interface mode.

```
R2(config -if)# exit
```

13. Enter the interface configuration mode

```
R2(config)#int xe23
```

14. Configure the IP address for the interface.

```
R2(config -if)# ip address 40.1.1.1/24
R2(config -if)# ipv6 address 4001::1/64
```

15. Include the interface in the router's ISIS 1 instance.

```
R2(config -if)# ip router isis 1
R2(config -if)# ipv6 router isis 1
```

16. Exit the interface mode.

```
R2(config -if)# exit
```

17. Set the routing process ID as 1.

```
R2(config)# router isis 1
```

18. Configure IS type as level 1.

```
R2(config-router)# is-type level-1
```

19. Configure wide metric style.

```
R2(config-router)# metric-style wide
```

20. Enable dynamic host name under ISIS process.

```
R2(config-router)# dynamic-hostname
```

21. Enable BFD in all the interfaces.

```
R2(config-router)# bfd all-interfaces
```

22. Configure Network Entity Title (NET).

```
R2(config-router)# net 49.0000.0000.0002.00
```

23. Commit the candidate configuration to the running configuration.

```
R2(config-router)# commit
```

R3

1. Enter configure mode followed by interface mode for loopback interface.

```
#configure terminal
R3(config)#int lo
```

2. Configure the IP address for the interface.

```
R3(config -if)# ip add 3.3.3.3/32 secondary
R3(config -if)# ipv6 address 3333::33/128
```

3. Include the interface in the router's ISIS 1 instance.

```
R3(config -if)# ip router isis 1
R3(config -if)# ipv6 router isis 1
```

4. Exit the interface mode.

```
R3(config -if)# exit
```

5. Enter the interface configuration mode.

```
R3(config)#int xe31/4
```

6. Configure the IP address for the interface.

```
R3(config -if)# ip address 40.1.1.2/24
R3(config -if)# ipv6 address 4001::2/64
```

7. Include the interface in the router's ISIS 1 instance.

```
R3(config -if)# ip router isis 1
R3(config -if)# ipv6 router isis 1
```

8. Exit interface mode

```
R3(config -if)# exit
```

9. Enter Interface configuration mode

```
R3(config)#int xe31/1
```

10. Configure the IP address of the interface

```
R3(config -if)# ip address 50.1.1.1/24
R3(config -if)# ipv6 address 5001::1/64
```

11. Include the interface in the router's ISIS 1 instance

```
R3(config -if)# ip router isis 1
R3(config -if)# ipv6 router isis 1
```

12. Exit interface mode

```
R3(config -if)# exit
```

13. Set the routing process ID as 1

```
R3(config)# router isis 1
```

14. Configure IS type as level-1

```
R3(config-router)# is-type level-1
```

15. Configure wide metric-style.

```
R3(config-router)# metric-style wide
```

16. Enable dynamic host name under ISIS process.

```
R3(config-router)# dynamic-hostname
```

17. Enable BFD on all the interfaces.

```
R3(config-router)# bfd all-interfaces
```

18. Configure Network Entity Title (NET).

```
R3(config-router)# net 49.0000.0000.0003.00
```

19. Commit the candidate configuration to the running configuration.

```
R3(config-router)# commit
```

R4**1. Enter configure mode followed by interface mode for loopback interface.**

```
#configure terminal  
R4(config)#int lo
```

2. Configure the IP address for the interface.

```
R4(config -if)# ip add 4.4.4.4/32 secondary  
R4(config -if)# ipv6 address 4444::22/128
```

3. Include the interface in the router's ISIS 1 instance.

```
R4(config -if)# ip router isis 1  
R4(config -if)# ipv6 router isis 1
```

4. Exit the interface mode.

```
R4(config -if)# exit
```

5. Enter the interface configuration mode.

```
R4(config)#int xe2
```

6. Configure the IP address of the interface.

```
R4(config -if)# ip address 20.1.1.2/24  
R4(config -if)# ipv6 address 2001::2/64
```

7. Include the interface in the router's ISIS 1 instance.

```
R4(config -if)# ipv6 router isis 1
```

8. Exit the interface mode.

```
R4(config -if)# exit
```

9. Enter the interface configuration mode.

```
R4(config)#int ge8
```

10. Configure the IP address for the interface.

```
R4(config -if)# ip address 50.1.1.2/24  
R4(config -if)# ipv6 address 5001::2/64
```

11. Include the interface in the router's ISIS 1 instance.

```
R4(config -if)# ip router isis 1  
R4(config -if)# ipv6 router isis 1
```

12. Exit the interface mode.

```
R4(config -if)# exit
```

13. Enter the interface configuration mode.

```
R4(config)#int xe1
```

14. Configure the IP address of the interface.

```
R4(config -if)# ip address 30.1.1.1/24  
R4(config -if)# ipv6 address 3001::1/64
```

15. Include the interface in the router's ISIS 1 instance

```
R4(config -if)# ip router isis 1
R4(config -if)# ipv6 router isis 1
```

16. Exit interface mode.

```
R4(config -if)# exit
```

17. Set the routing process ID as 1.

```
R4(config)# router isis 1
```

18. Configure IS type as level-1.

```
R4(config-router)# is-type level-1
```

19. Configure wide metric-style.

```
R4(config-router)# metric-style wide
```

20. Enable dynamic-hostname under ISIS process.

```
R4(config-router)# dynamic-hostname
```

21. Enable BFD on all the interfaces.

```
R4(config-router)# bfd all-interfaces
```

22. Configure Network Entity Title (NET).

```
R4(config-router)# net 49.0000.0000.0004.00
```

23. Commit the candidate configuration to the running configuration.

```
R4(config-router)# commit
```

R5**1. Enter configure mode followed by interface mode for loopback interface.**

```
#configure terminal
R5(config)#int lo
```

2. Configure the IP address of the interface.

```
R5(config -if)# ip add 5.5.5.5/32 secondary
R5(config -if)# ipv6 address 5555::55/128
```

3. Include the interface in the router's ISIS 1 instance.

```
R5(config -if)# ip router isis 1
R5(config -if)# ipv6 router isis 1
```

4. Exit interface mode.

```
R5(config -if)# exit
```

5. Enter interface configuration mode.

```
R5(config)#int ce22/1
```

6. Configure the IP address of the interface.

```
R5(config -if)# ip address 30.1.1.2/24
R5(config -if)# ipv6 address 3001::2/64
```

7. Include the interface in the router's ISIS 1 instance.

```
R5(config -if)# ip router isis 1
R5(config -if)# ipv6 router isis 1
```

8. Exit the interface mode.

```
R5(config-if)# exit
```

9. Set the routing process ID as 1.

```
R5(config)# router isis 1
```

10. Configure IS type as level-1.

```
R5(config-router)# is-type level-1
```

11. Configure wide metric-style.

```
R5(config-router)# metric-style wide
```

12. Enable dynamic host name under ISIS process.

```
R5(config-router)# dynamic-hostname
```

13. Enable BFD on all the interfaces.

```
R5(config-router)# bfd all-interfaces
```

14. Configure Network entity title (NET).

```
R5(config-router)# net 49.0000.0000.0005.00
```

15. Commit the candidate configuration to the running configuration.

```
R5(config-router)# commit
```

Once the configuration is done per router, follow the below mentioned steps to enable multi topology on all the routers.

Note: In the commands, modify the relevant router as R1, R2, R3, R4 or R5, depending on the router being configured.

1. Set the routing process ID as 1.

```
R1(config)# router isis 1
```

2. Configure metric-style wide.

```
R1(config-router)# metric-style wide
```

3. Configure address family IPv6.

```
R1(config-router)#address-family ipv6
```

4. Enable multi topology with level 1.

```
R1(config-router-af)#multi-topology level-1
```

5. Commit the candidate configuration to the running configuration.

```
R1(config-router-af)#commit
```

Validation for Multi Topology

```
R1#show clns neighbors
```

```
Total number of L1 adjacencies: 1
```

```
Total number of L2 adjacencies: 0
```

```
Total number of adjacencies: 1
```

```
Tag 1: VRF : default
```

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
R2	xe22	00e0.4b77.39fe	Up	19	L1	M-ISIS

```
R1#show clns is-neighbors detail
```

Tag 1: VRF : default

```

System Id      Interface  State  Type Priority  Circuit Id
R2             xe22     Up     L1   64        0000.0000.0001.02
  L1 Adjacency ID: 1
  L2 Adjacency ID: 2
  Uptime: 01:09:39
  Area Address(es): 49
  IP Address(es): 10.1.1.2
  IPv6 Address(es): fe80::2e0:4bff:fe77:39fe
  Topology: IPv4, IPv6
  Level-1 Protocols Supported: IPv4, IPv6
  Bidirectional Forwarding Detection is enabled
  Adjacency advertisement: Advertise
  
```

R1#show isis topology

Tag 1: VRF : default

IS-IS paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
R1	--			
R2	10	R2 xe22	00e0.4b77.39fe	
R3	20	R2 xe22	00e0.4b77.39fe	
R4	30	R2 xe22	00e0.4b77.39fe	
R5	40	R2 xe22	00e0.4b77.39fe	

R1#show ipv6 isis topology

Tag 1: VRF : default

IS-IS paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
R1	--			
R2	10	R2 xe22	00e0.4b77.39fe	
R3	20	R2 xe22	00e0.4b77.39fe	
R4	20	R2 xe22	00e0.4b77.39fe	
R5	30	R2 xe22	00e0.4b77.39fe	

R1#show ip route

```

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

IP Route Table for VRF "default"

```

C          1.1.1.1/32 is directly connected, lo, installed 01:55:53, last update
01:55:53 ago
i L1      2.2.2.2/32 [115/20] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
i L1      3.3.3.3/32 [115/30] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
i L1      4.4.4.4/32 [115/40] via 10.1.1.2, xe22, installed 00:09:50, last update
00:09:50 ago
i L1      5.5.5.5/32 [115/50] via 10.1.1.2, xe22, installed 00:09:50, last update
00:09:50 ago
C          10.1.1.0/24 is directly connected, xe22, installed 01:55:53, last update
01:55:53 ago
i L1      30.1.1.0/24 [115/40] via 10.1.1.2, xe22, installed 00:09:50, last update
00:09:50 ago
i L1      40.1.1.0/24 [115/20] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
i L1      50.1.1.0/24 [115/30] via 10.1.1.2, xe22, installed 01:09:50, last update
01:09:50 ago
C          127.0.0.0/8 is directly connected, lo, installed 01:57:14, last update
01:57:14 ago

```

Gateway of last resort is not set

R1#show ipv6 route

IPv6 Routing Table

Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,
N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
v - vrf leaked

Timers: Uptime

IP Route Table for VRF "default"

```

C          ::1/128 via ::, lo, installed 01:57:15, last update 01:57:15 ago
C          1001::/64 via ::, xe22, installed 01:32:33, last update 01:32:33 ago
C          1111::11/128 via ::, lo, installed 01:33:09, last update 01:33:09 ago
i L1      2001::/64 [115/20] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1      2222::22/128 [115/20] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
i L1      3001::/64 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1      3333::33/128 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
i L1      4001::/64 [115/20] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1      4444::44/128 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
i L1      5001::/64 [115/30] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51, last
update 00:09:51 ago
i L1      5555::55/128 [115/40] via fe80::2e0:4bff:fe77:39fe, xe22, installed 00:09:51,
last update 00:09:51 ago
C          fe80::/64 via ::, xe25, installed 01:56:18, last update 01:56:18 ago

```

```
R1#show isis spf-logs level-1-2
Tag 1: VRF : default
Level-1 spf logs:
Next SPF is not scheduled yet
SPF schedule delay min 0 secs 500 msecs
SPF schedule delay max 50 secs 0 msecs
SPF algorithm executed 12 times
SPF algorithm last executed 00:09:57.608 ago
```

```
R1#show isis database verbose
Tag 1: VRF : default
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R1.00-00      * 0x00000015  0x9E64        602           0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:       0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11
  Metric:      10          IS-Extended R1.02
  Metric:      10          IS (MT-IPv6) R1.02
  Metric:      10          IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:      10          IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:      10          IPv6 (MT-IPv6) 1111::11/128
  Metric:      10          IPv6 (MT-IPv6) 1001::/64
R1.02-00      * 0x0000000C  0x724E        602           0/0/0
  Metric:      0          IS-Extended R1.00
  Metric:      0          IS-Extended R2.00
R2.00-00      0x00000014  0x2A52        601           0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:       0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric:      10          IS-Extended R1.02
  Metric:      10          IS-Extended R3.03
  Metric:      10          IS (MT-IPv6) R1.02
  Metric:      10          IS (MT-IPv6) R3.03
  Metric:      10          IS (MT-IPv6) R4.04
  Metric:      10          IP-Extended 2.2.2.2/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:      10          IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:      10          IP-Extended 40.1.1.0/24
```

```

    Prefix Attribute Flags[0]: ELC Set
Metric: 10      IPv6 (MT-IPv6) 2222::22/128
Metric: 10      IPv6 (MT-IPv6) 1001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
Metric: 10      IPv6 (MT-IPv6) 2001::/64
R3.00-00      0x00000013  0x7FCC      601      0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:     R3
IP Address:   3.3.3.3
IPv6 Address: 3333::33
Metric: 10      IS-Extended R4.01
Metric: 10      IS-Extended R3.03
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IS (MT-IPv6) R3.03
Metric: 10      IP-Extended 3.3.3.3/32
Metric: 10      IP-Extended 50.1.1.0/24
Metric: 10      IP-Extended 40.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 3333::33/128
Metric: 10      IPv6 (MT-IPv6) 5001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
R3.03-00      0x0000000C  0x6D4E      601      0/0/0
Metric: 0      IS-Extended R3.00
Metric: 0      IS-Extended R2.00
R4.00-00      0x00000015  0x8C0D      601      0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:     R4
IP Address:   50.1.1.2
IPv6 Address: 5001::2
Metric: 10      IS-Extended R5.02
Metric: 10      IS-Extended R4.01
Metric: 10      IS (MT-IPv6) R5.02
Metric: 10      IS (MT-IPv6) R4.04
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
Metric: 10      IPv6 (MT-IPv6) 4444::44/128
Metric: 10      IPv6 (MT-IPv6) 3001::/64
Metric: 10      IPv6 (MT-IPv6) 2001::/64
Metric: 10      IPv6 (MT-IPv6) 5001::/64
R4.01-00      0x00000007  0x9A25      601      0/0/0
Metric: 0      IS-Extended R4.00
Metric: 0      IS-Extended R3.00

```

```

R4.04-00          0x0000000C  0x6751          601             0/0/0
  Metric: 0        IS-Extended R4.00
  Metric: 0        IS-Extended R2.00
R5.00-00          0x00000010  0xFA0F          601             0/0/0
  Area Address: 49
  Topology: IPv4 (0x0) IPv6 (0x2)
  NLPID: 0xCC 0x8E
  Hostname: R5
  IP Address: 5.5.5.5
  IPv6 Address: 5555::55
  Metric: 10       IS-Extended R5.02
  Metric: 10       IS (MT-IPv6) R5.02
  Metric: 10       IP-Extended 5.5.5.5/32
    Prefix Attribute Flags[0]: ELC Set
  Metric: 10       IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric: 10       IPv6 (MT-IPv6) 5555::55/128
  Metric: 10       IPv6 (MT-IPv6) 3001::/64
R5.02-00          0x00000007  0xA813          601             0/0/0
  Metric: 0        IS-Extended R5.00
  Metric: 0        IS-Extended R4.00

```

R1#show isis database detail

Tag 1: VRF : default

IS-IS Level-1 Link State Database:

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
R1.00-00	* 0x00000015	0x9E64	596	0/0/0
Area Address: 49				
Topology: IPv4 (0x0) IPv6 (0x2)				
NLPID: 0xCC 0x8E				
Hostname: R1				
IP Address: 1.1.1.1				
IPv6 Address: 1111::11				
Metric: 10 IS-Extended R1.02				
Metric: 10 IS (MT-IPv6) R1.02				
Metric: 10 IP-Extended 1.1.1.1/32				
Metric: 10 IP-Extended 10.1.1.0/24				
Metric: 10 IPv6 (MT-IPv6) 1111::11/128				
Metric: 10 IPv6 (MT-IPv6) 1001::/64				
R1.02-00	* 0x0000000C	0x724E	596	0/0/0
Metric: 0 IS-Extended R1.00				
Metric: 0 IS-Extended R2.00				
R2.00-00	0x00000014	0x2A52	595	0/0/0
Area Address: 49				
Topology: IPv4 (0x0) IPv6 (0x2)				
NLPID: 0xCC 0x8E				
Hostname: R2				
IP Address: 2.2.2.2				
IPv6 Address: 2222::22				

```

Metric: 10      IS-Extended R1.02
Metric: 10      IS-Extended R3.03
Metric: 10      IS (MT-IPv6) R1.02
Metric: 10      IS (MT-IPv6) R3.03
Metric: 10      IS (MT-IPv6) R4.04
Metric: 10      IP-Extended 2.2.2.2/32
Metric: 10      IP-Extended 10.1.1.0/24
Metric: 10      IP-Extended 40.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 2222::22/128
Metric: 10      IPv6 (MT-IPv6) 1001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
Metric: 10      IPv6 (MT-IPv6) 2001::/64
R3.00-00      0x00000013  0x7FCC      595      0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:     R3
IP Address:   3.3.3.3
IPv6 Address: 3333::33
Metric: 10      IS-Extended R4.01
Metric: 10      IS-Extended R3.03
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IS (MT-IPv6) R3.03
Metric: 10      IP-Extended 3.3.3.3/32
Metric: 10      IP-Extended 50.1.1.0/24
Metric: 10      IP-Extended 40.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 3333::33/128
Metric: 10      IPv6 (MT-IPv6) 5001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
R3.03-00      0x0000000C  0x6D4E      595      0/0/0
Metric: 0      IS-Extended R3.00
Metric: 0      IS-Extended R2.00
R4.00-00      0x00000015  0x8C0D      595      0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:     R4
IP Address:   50.1.1.2
IPv6 Address: 5001::2
Metric: 10      IS-Extended R5.02
Metric: 10      IS-Extended R4.01
Metric: 10      IS (MT-IPv6) R5.02
Metric: 10      IS (MT-IPv6) R4.04
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IP-Extended 50.1.1.0/24
Metric: 10      IP-Extended 4.4.4.4/32
Metric: 10      IP-Extended 30.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 4444::44/128
Metric: 10      IPv6 (MT-IPv6) 3001::/64
Metric: 10      IPv6 (MT-IPv6) 2001::/64

```



```

Metric: 10 IPv6 (MT-IPv6) 5001::/64
R4.01-00 0x00000007 0x9A25 595 0/0/0
Metric: 0 IS-Extended R4.00
Metric: 0 IS-Extended R3.00
R4.04-00 0x0000000C 0x6751 595 0/0/0
Metric: 0 IS-Extended R4.00
Metric: 0 IS-Extended R2.00
R5.00-00 0x00000010 0xFA0F 595 0/0/0
Area Address: 49
Topology: IPv4 (0x0) IPv6 (0x2)
NLPID: 0xCC 0x8E
Hostname: R5
IP Address: 5.5.5.5
IPv6 Address: 5555::55
Metric: 10 IS-Extended R5.02
Metric: 10 IS (MT-IPv6) R5.02
Metric: 10 IP-Extended 5.5.5.5/32
Metric: 10 IP-Extended 30.1.1.0/24
Metric: 10 IPv6 (MT-IPv6) 5555::55/128
Metric: 10 IPv6 (MT-IPv6) 3001::/64
R5.02-00 0x00000007 0xA813 595 0/0/0
Metric: 0 IS-Extended R5.00
Metric: 0 IS-Extended R4.00

```

R2:

R2#show clns neighbors

```

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1: VRF : default
System Id      Interface  SNPA                State  Holdtime  Type  Protocol
R1             xe12     e8c5.7a69.446f     Up     6         L1   M-ISIS
R3             xe23     903c.b3c5.ae9b     Up     6         L1   M-ISIS
R4             xe24     9819.2ccf.ede3     Up     9         L1   M-ISIS

```

R2#show clns is-neighbors detail

```

Tag 1: VRF : default
System Id      Interface  State  Type  Priority  Circuit Id
R1             xe12     Up     L1   64       0000.0000.0001.02
L1 Adjacency ID: 1
L2 Adjacency ID: 2
Uptime: 01:10:56
Area Address(es): 49
IP Address(es): 10.1.1.1
IPv6 Address(es): fe80::eac5:7aff:fe69:446f
Topology: IPv4, IPv6

```

Level-1 Protocols Supported: IPv4, IPv6
 Bidirectional Forwarding Detection is enabled
 Adjacency advertisement: Advertise

```
R3          xe23          Up      L1   64          0000.0000.0003.03
L1 Adjacency ID: 1
L2 Adjacency ID: 2
Uptime: 01:10:56
Area Address(es): 49
IP Address(es): 40.1.1.2
IPv6 Address(es): fe80::923c:b3ff:fec5:ae9b
Topology: IPv4, IPv6
Level-1 Protocols Supported: IPv4, IPv6
Bidirectional Forwarding Detection is enabled
Adjacency advertisement: Advertise
```

```
R4          xe24          Up      L1   64          0000.0000.0004.04
L1 Adjacency ID: 1
L2 Adjacency ID: 2
Uptime: 01:10:56
Area Address(es): 49
IPv6 Address(es): fe80::9a19:2cff:fecf:ede3
Topology: IPv6
Level-1 Protocols Supported: IPv4, IPv6
Bidirectional Forwarding Detection is enabled
Adjacency advertisement: Advertise
```

R2#show isis topology

```
Tag 1: VRF : default
IS-IS paths to level-1 routers
System Id      Metric      Next-Hop      Interface      SNPA
R1             10         R1            xe12           e8c5.7a69.446f
R2             --
R3             10         R3            xe23           903c.b3c5.ae9b
R4             20         R3            xe23           903c.b3c5.ae9b
R5             30         R3            xe23           903c.b3c5.ae9b
```

R2#show ipv6 isis topology

```
Tag 1: VRF : default
IS-IS paths to level-1 routers
System Id      Metric      Next-Hop      Interface      SNPA
R1             10         R1            xe12           e8c5.7a69.446f
R2             --
R3             10         R3            xe23           903c.b3c5.ae9b
R4             10         R4            xe24           9819.2ccf.ede3
R5             20         R4            xe24           9819.2ccf.ede3
```

```
R2#show ip route
```

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

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

```
i L1      1.1.1.1/32 [115/20] via 10.1.1.1, xe12, installed 01:11:03, last update
01:11:03 ago
C         2.2.2.2/32 is directly connected, lo, installed 01:59:20, last update
01:59:20 ago
i L1      3.3.3.3/32 [115/20] via 40.1.1.2, xe23, installed 01:11:03, last update
01:11:03 ago
i L1      4.4.4.4/32 [115/30] via 40.1.1.2, xe23, installed 00:11:03, last update
00:11:03 ago
i L1      5.5.5.5/32 [115/40] via 40.1.1.2, xe23, installed 00:11:03, last update
00:11:03 ago
C         10.1.1.0/24 is directly connected, xe12, installed 01:57:30, last update
01:57:30 ago
C         20.1.1.0/24 is directly connected, xe24, installed 01:59:19, last update
01:59:19 ago
i L1      30.1.1.0/24 [115/30] via 40.1.1.2, xe23, installed 00:11:03, last update
00:11:03 ago
C         40.1.1.0/24 is directly connected, xe23, installed 01:59:19, last update
01:59:19 ago
i L1      50.1.1.0/24 [115/20] via 40.1.1.2, xe23, installed 01:11:03, last update
01:11:03 ago
C         127.0.0.0/8 is directly connected, lo, installed 02:20:04, last update
02:20:04 ago
```

```
Gateway of last resort is not set
```

```
R2#show ipv6 route
```

```
IPv6 Routing Table
```

```
Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
       O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
       E2 - OSPF external type 2, E - EVPN  N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
       P - SRV6-POLICY,
       v - vrf leaked
```

```
Timers: Uptime
```

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

```
C        ::1/128 via ::, lo, installed 02:20:05, last update 02:20:05 ago
C        1001::/64 via ::, xe12, installed 01:32:42, last update 01:32:42 ago
i L1     1111::11/128 [115/20] via fe80::eac5:7aff:fe69:446f, xe12, installed 00:11:04,
last update 00:11:04 ago
C        2001::/64 via ::, xe24, installed 01:59:20, last update 01:59:20 ago
```

```

C      2222::22/128 via ::, lo, installed 01:33:21, last update 01:33:21 ago
i L1   3001::/64 [115/20] via fe80::9a19:2cff:fe3f:ede3, xe24, installed 00:11:04, last
update 00:11:04 ago
i L1   3333::33/128 [115/20] via fe80::923c:b3ff:fe35:ae9b, xe23, installed 01:11:04,
last update 01:11:04 ago
C      4001::/64 via ::, xe23, installed 01:24:52, last update 01:24:52 ago
i L1   4444::44/128 [115/20] via fe80::9a19:2cff:fe3f:ede3, xe24, installed 00:11:04,
last update 00:11:04 ago
i L1   5001::/64 [115/20] via fe80::923c:b3ff:fe35:ae9b, xe23, installed 01:11:04, last
update 00:11:04 ago
      [115/20] via fe80::9a19:2cff:fe3f:ede3, xe24
i L1   5555::55/128 [115/30] via fe80::9a19:2cff:fe3f:ede3, xe24, installed 00:11:04,
last update 00:11:04 ago
C      fe80::/64 via ::, xe12, installed 01:57:31, last update 01:57:31 ago

```

```

R2#show isis spf-logs level-1-2
Tag 1: VRF : default
Level-1 spf logs:
Next SPF is not scheduled yet
SPF schedule delay min 0 secs 500 msec
SPF schedule delay max 50 secs 0 msec
SPF algorithm executed 12 times
SPF algorithm last executed 00:11:11.544 ago

```

```

R2#show isis database verbose
Tag 1: VRF : default
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R1.00-00       0x00000015   0x9E64        527            0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:       0xCC 0x8E
  Hostname:    R1
  IP Address:  1.1.1.1
  IPv6 Address: 1111::11
  Metric:     10          IS-Extended R1.02
  Metric:     10          IS (MT-IPv6) R1.02
  Metric:     10          IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:     10          IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:     10          IPv6 (MT-IPv6) 1111::11/128
  Metric:     10          IPv6 (MT-IPv6) 1001::/64
R1.02-00       0x0000000C   0x724E        527            0/0/0
  Metric:     0          IS-Extended R1.00
  Metric:     0          IS-Extended R2.00
R2.00-00       * 0x00000014  0x2A52        528            0/0/0
  Area Address: 49

```

```

Topology:      IPv4 (0x0) IPv6 (0x2)
NLPID:        0xCC 0x8E
Hostname:     R2
IP Address:   2.2.2.2
IPv6 Address: 2222::22
Metric:      10          IS-Extended R1.02
Metric:      10          IS-Extended R3.03
Metric:      10          IS (MT-IPv6) R1.02
Metric:      10          IS (MT-IPv6) R3.03
Metric:      10          IS (MT-IPv6) R4.04
Metric:      10          IP-Extended 2.2.2.2/32
  Prefix Attribute Flags[0]: ELC Set
Metric:      10          IP-Extended 10.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric:      10          IP-Extended 40.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric:      10          IPv6 (MT-IPv6) 2222::22/128
Metric:      10          IPv6 (MT-IPv6) 1001::/64
Metric:      10          IPv6 (MT-IPv6) 4001::/64
Metric:      10          IPv6 (MT-IPv6) 2001::/64
R3.00-00      0x00000013   0x7FCC           527           0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:    R3
IP Address:  3.3.3.3
IPv6 Address: 3333::33
Metric:     10          IS-Extended R4.01
Metric:     10          IS-Extended R3.03
Metric:     10          IS (MT-IPv6) R4.01
Metric:     10          IS (MT-IPv6) R3.03
Metric:     10          IP-Extended 3.3.3.3/32
Metric:     10          IP-Extended 50.1.1.0/24
Metric:     10          IP-Extended 40.1.1.0/24
Metric:     10          IPv6 (MT-IPv6) 3333::33/128
Metric:     10          IPv6 (MT-IPv6) 5001::/64
Metric:     10          IPv6 (MT-IPv6) 4001::/64
R3.03-00      0x0000000C   0x6D4E           527           0/0/0
Metric:      0          IS-Extended R3.00
Metric:      0          IS-Extended R2.00
R4.00-00      0x00000015   0x8C0D           527           0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:    R4
IP Address:  50.1.1.2
IPv6 Address: 5001::2
Metric:     10          IS-Extended R5.02
Metric:     10          IS-Extended R4.01
Metric:     10          IS (MT-IPv6) R5.02

```

```

Metric: 10          IS (MT-IPv6) R4.04
Metric: 10          IS (MT-IPv6) R4.01
Metric: 10          IP-Extended 50.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 4.4.4.4/32
    Prefix Attribute Flags[0]: ELC Set
Metric: 10          IP-Extended 30.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
Metric: 10          IPv6 (MT-IPv6) 4444::44/128
Metric: 10          IPv6 (MT-IPv6) 3001::/64
Metric: 10          IPv6 (MT-IPv6) 2001::/64
Metric: 10          IPv6 (MT-IPv6) 5001::/64
R4.01-00            0x00000007  0x9A25          527          0/0/0
    Metric: 0          IS-Extended R4.00
    Metric: 0          IS-Extended R3.00
R4.04-00            0x0000000C  0x6751          527          0/0/0
    Metric: 0          IS-Extended R4.00
    Metric: 0          IS-Extended R2.00
R5.00-00            0x00000010  0xFA0F          527          0/0/0
    Area Address: 49
    Topology: IPv4 (0x0) IPv6 (0x2)
    NLPID: 0xCC 0x8E
    Hostname: R5
    IP Address: 5.5.5.5
    IPv6 Address: 5555::55
    Metric: 10          IS-Extended R5.02
    Metric: 10          IS (MT-IPv6) R5.02
    Metric: 10          IP-Extended 5.5.5.5/32
        Prefix Attribute Flags[0]: ELC Set
    Metric: 10          IP-Extended 30.1.1.0/24
        Prefix Attribute Flags[0]: ELC Set
    Metric: 10          IPv6 (MT-IPv6) 5555::55/128
    Metric: 10          IPv6 (MT-IPv6) 3001::/64
R5.02-00            0x00000007  0xA813          527          0/0/0
    Metric: 0          IS-Extended R5.00
    Metric: 0          IS-Extended R4.00

```

R2#show isis database detail

Tag 1: VRF : default

IS-IS Level-1 Link State Database:

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
R1.00-00	0x00000015	0x9E64	520	0/0/0

Area Address: 49

Topology: IPv4 (0x0) IPv6 (0x2)

NLPID: 0xCC 0x8E

Hostname: R1

IP Address: 1.1.1.1

IPv6 Address: 1111::11

Metric: 10 IS-Extended R1.02

Metric: 10 IS (MT-IPv6) R1.02

```

Metric: 10      IP-Extended 1.1.1.1/32
Metric: 10      IP-Extended 10.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 1111::11/128
Metric: 10      IPv6 (MT-IPv6) 1001::/64
R1.02-00      0x0000000C  0x724E      520      0/0/0
Metric: 0      IS-Extended R1.00
Metric: 0      IS-Extended R2.00
R2.00-00      * 0x00000014  0x2A52      521      0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:    R2
IP Address:  2.2.2.2
IPv6 Address: 2222::22
Metric: 10      IS-Extended R1.02
Metric: 10      IS-Extended R3.03
Metric: 10      IS (MT-IPv6) R1.02
Metric: 10      IS (MT-IPv6) R3.03
Metric: 10      IS (MT-IPv6) R4.04
Metric: 10      IP-Extended 2.2.2.2/32
Metric: 10      IP-Extended 10.1.1.0/24
Metric: 10      IP-Extended 40.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 2222::22/128
Metric: 10      IPv6 (MT-IPv6) 1001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
Metric: 10      IPv6 (MT-IPv6) 2001::/64
R3.00-00      0x00000013  0x7FCC      520      0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:    R3
IP Address:  3.3.3.3
IPv6 Address: 3333::33
Metric: 10      IS-Extended R4.01
Metric: 10      IS-Extended R3.03
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IS (MT-IPv6) R3.03
Metric: 10      IP-Extended 3.3.3.3/32
Metric: 10      IP-Extended 50.1.1.0/24
Metric: 10      IP-Extended 40.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 3333::33/128
Metric: 10      IPv6 (MT-IPv6) 5001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
R3.03-00      0x0000000C  0x6D4E      520      0/0/0
Metric: 0      IS-Extended R3.00
Metric: 0      IS-Extended R2.00
R4.00-00      0x00000015  0x8C0D      520      0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E

```

```

Hostname:      R4
IP Address:    50.1.1.2
IPv6 Address:  5001::2
Metric: 10    IS-Extended R5.02
Metric: 10    IS-Extended R4.01
Metric: 10    IS (MT-IPv6) R5.02
Metric: 10    IS (MT-IPv6) R4.04
Metric: 10    IS (MT-IPv6) R4.01
Metric: 10    IP-Extended 50.1.1.0/24
Metric: 10    IP-Extended 4.4.4.4/32
Metric: 10    IP-Extended 30.1.1.0/24
Metric: 10    IPv6 (MT-IPv6) 4444::44/128
Metric: 10    IPv6 (MT-IPv6) 3001::/64
Metric: 10    IPv6 (MT-IPv6) 2001::/64
Metric: 10    IPv6 (MT-IPv6) 5001::/64
R4.01-00      0x00000007  0x9A25      520      0/0/0
  Metric: 0    IS-Extended R4.00
  Metric: 0    IS-Extended R3.00
R4.04-00      0x0000000C  0x6751      520      0/0/0
  Metric: 0    IS-Extended R4.00
  Metric: 0    IS-Extended R2.00
R5.00-00      0x00000010  0xFA0F      520      0/0/0
  Area Address: 49
  Topology:    IPv4 (0x0) IPv6 (0x2)
  NLPID:      0xCC 0x8E
  Hostname:    R5
  IP Address:  5.5.5.5
  IPv6 Address: 5555::55
  Metric: 10    IS-Extended R5.02
  Metric: 10    IS (MT-IPv6) R5.02
  Metric: 10    IP-Extended 5.5.5.5/32
  Metric: 10    IP-Extended 30.1.1.0/24
  Metric: 10    IPv6 (MT-IPv6) 5555::55/128
  Metric: 10    IPv6 (MT-IPv6) 3001::/64
R5.02-00      0x00000007  0xA813      520      0/0/0
  Metric: 0    IS-Extended R5.00
  Metric: 0    IS-Extended R4.00

```

R3:

R3#show clns neighbors

```

Total number of L1 adjacencies: 2
Total number of L2 adjacencies: 0
Total number of adjacencies: 2
Tag 1: VRF : default
System Id      Interface  SNPA      State  Holdtime  Type Protocol

```


R4	xe31/1	9819.2ccf.ede9	Up	9	L1	M-ISIS
R2	xe31/4	00e0.4b77.3a09	Up	27	L1	M-ISIS

R3#show clns is-neighbors detail

Tag 1: VRF : default

System Id	Interface	State	Type	Priority	Circuit Id
R4	xe31/1	Up	L1	64	0000.0000.0004.01

L1 Adjacency ID: 1

L2 Adjacency ID: 2

Uptime: 01:11:42

Area Address(es): 49

IP Address(es): 50.1.1.2

IPv6 Address(es): fe80::9a19:2cff:fe9f:ede9

Topology: IPv4, IPv6

Level-1 Protocols Supported: IPv4, IPv6

Bidirectional Forwarding Detection is enabled

Adjacency advertisement: Advertise

R2	xe31/4	Up	L1	64	0000.0000.0003.03
----	--------	----	----	----	-------------------

L1 Adjacency ID: 1

L2 Adjacency ID: 2

Uptime: 01:11:42

Area Address(es): 49

IP Address(es): 40.1.1.1

IPv6 Address(es): fe80::2e0:4bff:fe77:3a09

Topology: IPv4, IPv6

Level-1 Protocols Supported: IPv4, IPv6

Bidirectional Forwarding Detection is enabled

Adjacency advertisement: Advertise

R3#show isis topology

Tag 1: VRF : default

IS-IS paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
R1	20	R2	xe31/4	00e0.4b77.3a09
R2	10	R2	xe31/4	00e0.4b77.3a09
R3	--			
R4	10	R4	xe31/1	9819.2ccf.ede9
R5	20	R4	xe31/1	9819.2ccf.ede9

R3#show ipv6 isis topology

Tag 1: VRF : default

IS-IS paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
R1	20	R2	xe31/4	00e0.4b77.3a09

R2	10	R2	xe31/4	00e0.4b77.3a09
R3	--			
R4	10	R4	xe31/1	9819.2ccf.ede9
R5	20	R4	xe31/1	9819.2ccf.ede9

R3#show ip route

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

IP Route Table for VRF "default"

```
i L1      1.1.1.1/32 [115/30] via 40.1.1.1, xe31/4, installed 01:11:53, last update
01:11:53 ago
i L1      2.2.2.2/32 [115/20] via 40.1.1.1, xe31/4, installed 01:11:53, last update
01:11:53 ago
C         3.3.3.3/32 is directly connected, lo, installed 02:00:27, last update
02:00:27 ago
i L1      4.4.4.4/32 [115/20] via 50.1.1.2, xe31/1, installed 01:11:53, last update
01:11:53 ago
i L1      5.5.5.5/32 [115/30] via 50.1.1.2, xe31/1, installed 01:11:53, last update
01:11:53 ago
i L1      10.1.1.0/24 [115/20] via 40.1.1.1, xe31/4, installed 01:11:53, last update
01:11:53 ago
i L1      30.1.1.0/24 [115/20] via 50.1.1.2, xe31/1, installed 01:11:53, last update
01:11:53 ago
C         40.1.1.0/24 is directly connected, xe31/4, installed 02:00:09, last update
02:00:09 ago
C         50.1.1.0/24 is directly connected, xe31/1, installed 02:00:26, last update
02:00:26 ago
C         127.0.0.0/8 is directly connected, lo, installed 02:18:52, last update
02:18:52 ago
```

Gateway of last resort is not set

R3#show ipv6 route

IPv6 Routing Table

Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
 O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
 E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,
 N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
 v - vrf leaked

Timers: Uptime

IP Route Table for VRF "default"

```
C         ::1/128 via ::, lo, installed 02:18:53, last update 02:18:53 ago
i L1      1001::/64 [115/20] via fe80::2e0:4bff:fe77:3a09, xe31/4, installed 00:11:54,
last update 00:11:54 ago
```

```

i L1 1111::11/128 [115/30] via fe80::2e0:4bff:fe77:3a09, xe31/4, installed 00:11:54,
last update 00:11:54 ago
i L1 2001::/64 [115/20] via fe80::9a19:2cff:fe77:3a09, xe31/1, installed 00:11:54,
last update 00:11:54 ago
      [115/20] via fe80::2e0:4bff:fe77:3a09, xe31/4
i L1 2222::22/128 [115/20] via fe80::2e0:4bff:fe77:3a09, xe31/4, installed 00:11:54,
last update 00:11:54 ago
i L1 3001::/64 [115/20] via fe80::9a19:2cff:fe77:3a09, xe31/1, installed 00:11:54,
last update 00:11:54 ago
C 3333::33/128 via ::, lo, installed 01:31:50, last update 01:31:50 ago
C 4001::/64 via ::, xe31/4, installed 01:30:10, last update 01:30:10 ago
i L1 4444::44/128 [115/20] via fe80::9a19:2cff:fe77:3a09, xe31/1, installed 00:11:54,
last update 00:11:54 ago
C 5001::/64 via ::, xe31/1, installed 01:29:43, last update 01:29:43 ago
i L1 5555::55/128 [115/30] via fe80::9a19:2cff:fe77:3a09, xe31/1, installed 00:11:54,
last update 00:11:54 ago
C fe80::/64 via ::, xe31/4, installed 02:00:10, last update 02:00:10 ago

```

```

R3#show isis spf-logs level-1-2
Tag 1: VRF : default
Level-1 spf logs:
Next SPF is not scheduled yet
SPF schedule delay min 0 secs 500 msec
SPF schedule delay max 50 secs 0 msec
SPF algorithm executed 12 times
SPF algorithm last executed 00:12:00.519 ago

```

```

R3#show isis database verbose
Tag 1: VRF : default
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R1.00-00      0x00000015  0x9E64        478           0/0/0
  Area Address: 49
  Topology:     IPv4 (0x0) IPv6 (0x2)
  NLPID:       0xCC 0x8E
  Hostname:    R1
  IP Address:  1.1.1.1
  IPv6 Address: 1111::11
  Metric:     10          IS-Extended R1.02
  Metric:     10          IS (MT-IPv6) R1.02
  Metric:     10          IP-Extended 1.1.1.1/32
    Prefix Attribute Flags[0]: ELC Set
  Metric:     10          IP-Extended 10.1.1.0/24
    Prefix Attribute Flags[0]: ELC Set
  Metric:     10          IPv6 (MT-IPv6) 1111::11/128
  Metric:     10          IPv6 (MT-IPv6) 1001::/64
R1.02-00      0x0000000C  0x724E        478           0/0/0
  Metric:     0          IS-Extended R1.00
  Metric:     0          IS-Extended R2.00

```

```

R2.00-00          0x00000014   0x2A52          478          0/0/0
  Area Address: 49
  Topology:      IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R2
  IP Address:   2.2.2.2
  IPv6 Address: 2222::22
  Metric: 10    IS-Extended R1.02
  Metric: 10    IS-Extended R3.03
  Metric: 10    IS (MT-IPv6) R1.02
  Metric: 10    IS (MT-IPv6) R3.03
  Metric: 10    IS (MT-IPv6) R4.04
  Metric: 10    IP-Extended 2.2.2.2/32
  Prefix Attribute Flags[0]: ELC Set
  Metric: 10    IP-Extended 10.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
  Metric: 10    IP-Extended 40.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
  Metric: 10    IPv6 (MT-IPv6) 2222::22/128
  Metric: 10    IPv6 (MT-IPv6) 1001::/64
  Metric: 10    IPv6 (MT-IPv6) 4001::/64
  Metric: 10    IPv6 (MT-IPv6) 2001::/64
R3.00-00          * 0x00000013   0x7FCC          479          0/0/0
  Area Address: 49
  Topology:      IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R3
  IP Address:   3.3.3.3
  IPv6 Address: 3333::33
  Metric: 10    IS-Extended R4.01
  Metric: 10    IS-Extended R3.03
  Metric: 10    IS (MT-IPv6) R4.01
  Metric: 10    IS (MT-IPv6) R3.03
  Metric: 10    IP-Extended 3.3.3.3/32
  Metric: 10    IP-Extended 50.1.1.0/24
  Metric: 10    IP-Extended 40.1.1.0/24
  Metric: 10    IPv6 (MT-IPv6) 3333::33/128
  Metric: 10    IPv6 (MT-IPv6) 5001::/64
  Metric: 10    IPv6 (MT-IPv6) 4001::/64
R3.03-00          * 0x0000000C   0x6D4E          479          0/0/0
  Metric: 0     IS-Extended R3.00
  Metric: 0     IS-Extended R2.00
R4.00-00          0x00000015   0x8C0D          478          0/0/0
  Area Address: 49
  Topology:      IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R4
  IP Address:   50.1.1.2
  IPv6 Address: 5001::2
  Metric: 10    IS-Extended R5.02

```

```

Metric: 10      IS-Extended R4.01
Metric: 10      IS (MT-IPv6) R5.02
Metric: 10      IS (MT-IPv6) R4.04
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IP-Extended 50.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 4.4.4.4/32
  Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 30.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric: 10      IPv6 (MT-IPv6) 4444::44/128
Metric: 10      IPv6 (MT-IPv6) 3001::/64
Metric: 10      IPv6 (MT-IPv6) 2001::/64
Metric: 10      IPv6 (MT-IPv6) 5001::/64
R4.01-00      0x00000007  0x9A25      478      0/0/0
  Metric: 0      IS-Extended R4.00
  Metric: 0      IS-Extended R3.00
R4.04-00      0x0000000C  0x6751      478      0/0/0
  Metric: 0      IS-Extended R4.00
  Metric: 0      IS-Extended R2.00
R5.00-00      0x00000010  0xFA0F      478      0/0/0
  Area Address: 49
  Topology:      IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R5
  IP Address:   5.5.5.5
  IPv6 Address: 5555::55
  Metric: 10      IS-Extended R5.02
  Metric: 10      IS (MT-IPv6) R5.02
  Metric: 10      IP-Extended 5.5.5.5/32
  Prefix Attribute Flags[0]: ELC Set
  Metric: 10      IP-Extended 30.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
  Metric: 10      IPv6 (MT-IPv6) 5555::55/128
  Metric: 10      IPv6 (MT-IPv6) 3001::/64
R5.02-00      0x00000007  0xA813      478      0/0/0
  Metric: 0      IS-Extended R5.00
  Metric: 0      IS-Extended R4.00

```

R3#show isis database detail

Tag 1: VRF : default

IS-IS Level-1 Link State Database:

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
R1.00-00	0x00000015	0x9E64	471	0/0/0

```

  Area Address: 49
  Topology:      IPv4 (0x0) IPv6 (0x2)
  NLPID:        0xCC 0x8E
  Hostname:     R1
  IP Address:   1.1.1.1
  IPv6 Address: 1111::11

```

```

Metric: 10      IS-Extended R1.02
Metric: 10      IS (MT-IPv6) R1.02
Metric: 10      IP-Extended 1.1.1.1/32
Metric: 10      IP-Extended 10.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 1111::11/128
Metric: 10      IPv6 (MT-IPv6) 1001::/64
R1.02-00        0x0000000C    0x724E          471          0/0/0
Metric: 0      IS-Extended R1.00
Metric: 0      IS-Extended R2.00
R2.00-00        0x00000014    0x2A52          471          0/0/0
Area Address: 49
Topology:      IPv4 (0x0) IPv6 (0x2)
NLPID:        0xCC 0x8E
Hostname:      R2
IP Address:    2.2.2.2
IPv6 Address: 2222::22
Metric: 10      IS-Extended R1.02
Metric: 10      IS-Extended R3.03
Metric: 10      IS (MT-IPv6) R1.02
Metric: 10      IS (MT-IPv6) R3.03
Metric: 10      IS (MT-IPv6) R4.04
Metric: 10      IP-Extended 2.2.2.2/32
Metric: 10      IP-Extended 10.1.1.0/24
Metric: 10      IP-Extended 40.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 2222::22/128
Metric: 10      IPv6 (MT-IPv6) 1001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
Metric: 10      IPv6 (MT-IPv6) 2001::/64
R3.00-00        * 0x00000013    0x7FCC          472          0/0/0
Area Address: 49
Topology:      IPv4 (0x0) IPv6 (0x2)
NLPID:        0xCC 0x8E
Hostname:      R3
IP Address:    3.3.3.3
IPv6 Address: 3333::33
Metric: 10      IS-Extended R4.01
Metric: 10      IS-Extended R3.03
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IS (MT-IPv6) R3.03
Metric: 10      IP-Extended 3.3.3.3/32
Metric: 10      IP-Extended 50.1.1.0/24
Metric: 10      IP-Extended 40.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 3333::33/128
Metric: 10      IPv6 (MT-IPv6) 5001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
R3.03-00        * 0x0000000C    0x6D4E          472          0/0/0
Metric: 0      IS-Extended R3.00
Metric: 0      IS-Extended R2.00
R4.00-00        0x00000015    0x8C0D          471          0/0/0
Area Address: 49

```

```

Topology:      IPv4 (0x0) IPv6 (0x2)
NLPID:        0xCC 0x8E
Hostname:     R4
IP Address:   50.1.1.2
IPv6 Address: 5001::2
Metric:      10          IS-Extended R5.02
Metric:      10          IS-Extended R4.01
Metric:      10          IS (MT-IPv6) R5.02
Metric:      10          IS (MT-IPv6) R4.04
Metric:      10          IS (MT-IPv6) R4.01
Metric:      10          IP-Extended 50.1.1.0/24
Metric:      10          IP-Extended 4.4.4.4/32
Metric:      10          IP-Extended 30.1.1.0/24
Metric:      10          IPv6 (MT-IPv6) 4444::44/128
Metric:      10          IPv6 (MT-IPv6) 3001::/64
Metric:      10          IPv6 (MT-IPv6) 2001::/64
Metric:      10          IPv6 (MT-IPv6) 5001::/64
R4.01-00      0x00000007  0x9A25          471          0/0/0
  Metric:     0          IS-Extended R4.00
  Metric:     0          IS-Extended R3.00
R4.04-00      0x0000000C  0x6751          471          0/0/0
  Metric:     0          IS-Extended R4.00
  Metric:     0          IS-Extended R2.00
R5.00-00      0x00000010  0xFA0F          471          0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:    R5
IP Address:  5.5.5.5
IPv6 Address: 5555::55
Metric:     10          IS-Extended R5.02
Metric:     10          IS (MT-IPv6) R5.02
Metric:     10          IP-Extended 5.5.5.5/32
Metric:     10          IP-Extended 30.1.1.0/24
Metric:     10          IPv6 (MT-IPv6) 5555::55/128
Metric:     10          IPv6 (MT-IPv6) 3001::/64
R5.02-00      0x00000007  0xA813          471          0/0/0
  Metric:     0          IS-Extended R5.00
  Metric:     0          IS-Extended R4.00

```

R4:

R4#show clns neighbors

```

Total number of L1 adjacencies: 3
Total number of L2 adjacencies: 0
Total number of adjacencies: 3
Tag 1: VRF : default

```

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
R5	xe1	e001.a6aa.0f23	Up	6	L1	M-ISIS
R2	xe2	00e0.4b77.3a0a	Up	22	L1	M-ISIS
R3	ge8	903c.b3c5.ae98	Up	22	L1	M-ISIS

R4#show clns is-neighbors detail

Tag 1: VRF : default

System Id	Interface	State	Type	Priority	Circuit Id
R5	xe1	Up	L1	64	0000.0000.0005.02

L1 Adjacency ID: 1
 L2 Adjacency ID: 2
 Uptime: 01:12:38
 Area Address(es): 49
 IP Address(es): 30.1.1.2
 IPv6 Address(es): fe80::e201:a6ff:feaa:f23
 Topology: IPv4, IPv6
 Level-1 Protocols Supported: IPv4, IPv6
 Bidirectional Forwarding Detection is enabled
 Adjacency advertisement: Advertise

R2	xe2	Up	L1	64	0000.0000.0004.04
----	-----	----	----	----	-------------------

L1 Adjacency ID: 1
 L2 Adjacency ID: 2
 Uptime: 01:12:37
 Area Address(es): 49
 IPv6 Address(es): fe80::2e0:4bff:fe77:3a0a
 Topology: IPv6
 Level-1 Protocols Supported: IPv4, IPv6
 Bidirectional Forwarding Detection is enabled
 Adjacency advertisement: Advertise

R3	ge8	Up	L1	64	0000.0000.0004.01
----	-----	----	----	----	-------------------

L1 Adjacency ID: 1
 L2 Adjacency ID: 2
 Uptime: 01:12:38
 Area Address(es): 49
 IP Address(es): 50.1.1.1
 IPv6 Address(es): fe80::923c:b3ff:fec5:ae98
 Topology: IPv4, IPv6
 Level-1 Protocols Supported: IPv4, IPv6
 Bidirectional Forwarding Detection is enabled
 Adjacency advertisement: Advertise

R4#show isis topology

Tag 1: VRF : default

IS-IS paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
-----------	--------	----------	-----------	------

R1	30	R3	ge8	903c.b3c5.ae98
R2	20	R3	ge8	903c.b3c5.ae98
R3	10	R3	ge8	903c.b3c5.ae98
R4	--			
R5	10	R5	xe1	e001.a6aa.0f23

R4#show ipv6 isis topology

Tag 1: VRF : default

IS-IS paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
R1	20	R2 xe2	00e0.4b77.3a0a	
R2	10	R2 xe2	00e0.4b77.3a0a	
R3	10	R3 ge8	903c.b3c5.ae98	
R4	--			
R5	10	R5 xe1	e001.a6aa.0f23	

R4#show ip route

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

O - OSPF, IA - OSPF inter area

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

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

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

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

v - vrf leaked

* - candidate default

IP Route Table for VRF "default"

i L1	1.1.1.1/32 [115/40]	via 50.1.1.1, ge8, installed 00:12:48, last update 00:12:48 ago
i L1	2.2.2.2/32 [115/30]	via 50.1.1.1, ge8, installed 00:12:48, last update 00:12:48 ago
i L1	3.3.3.3/32 [115/20]	via 50.1.1.1, ge8, installed 01:01:13, last update 01:01:13 ago
C	4.4.4.4/32	is directly connected, lo, installed 02:01:55, last update 02:01:55 ago
i L1	5.5.5.5/32 [115/20]	via 30.1.1.2, xe1, installed 01:12:47, last update 01:12:47 ago
i L1	10.1.1.0/24 [115/30]	via 50.1.1.1, ge8, installed 00:12:48, last update 00:12:48 ago
C	20.1.1.0/24	is directly connected, xe2, installed 02:01:04, last update 02:01:04 ago
C	30.1.1.0/24	is directly connected, xe1, installed 02:01:55, last update 02:01:55 ago
i L1	40.1.1.0/24 [115/20]	via 50.1.1.1, ge8, installed 01:01:13, last update 01:01:13 ago
C	50.1.1.0/24	is directly connected, ge8, installed 02:01:22, last update 02:01:22 ago
C	127.0.0.0/8	is directly connected, lo, installed 02:20:17, last update 02:20:17 ago

Gateway of last resort is not set

R4#show ipv6 route

IPv6 Routing Table

Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,
 O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,
 E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,
 N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,
 v - vrf leaked

Timers: Uptime

IP Route Table for VRF "default"

```

C      ::1/128 via ::, lo, installed 02:20:18, last update 02:20:18 ago
i L1   1001::/64 [115/20] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 00:12:48, last
update 00:12:48 ago
i L1   1111::11/128 [115/30] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 00:12:48,
last update 00:12:48 ago
C      2001::/64 via ::, xe2, installed 02:01:05, last update 02:01:05 ago
i L1   2222::22/128 [115/20] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 00:12:48,
last update 00:12:48 ago
C      3001::/64 via ::, xe1, installed 01:33:20, last update 01:33:20 ago
i L1   3333::33/128 [115/20] via fe80::923c:b3ff:fec5:ae98, ge8, installed 01:01:14,
last update 01:01:14 ago
i L1   4001::/64 [115/20] via fe80::2e0:4bff:fe77:3a0a, xe2, installed 01:04:04, last
update 00:12:48 ago
      [115/20] via fe80::923c:b3ff:fec5:ae98, ge8
C      4444::44/128 via ::, lo, installed 01:33:04, last update 01:33:04 ago
C      5001::/64 via ::, ge8, installed 01:29:27, last update 01:29:27 ago
i L1   5555::55/128 [115/20] via fe80::e201:a6ff:feaa:f23, xe1, installed 00:12:48,
last update 00:12:48 ago
C      fe80::/64 via ::, xe2, installed 02:01:05, last update 02:01:05 ago
    
```

R4#show isis spf-logs level-1-2

Tag 1: VRF : default

Level-1 spf logs:

```

Next SPF is not scheduled yet
SPF schedule delay min 0 secs 500 msec
SPF schedule delay max 50 secs 0 msec
SPF algorithm executed 12 times
SPF algorithm last executed 00:12:55.361 ago
    
```

R4#show isis database verbose

Tag 1: VRF : default

IS-IS Level-1 Link State Database:

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
R1.00-00	0x00000015	0x9E64	423	0/0/0
Area Address: 49				
Topology: IPv4 (0x0) IPv6 (0x2)				

```

NLPID:      0xCC 0x8E
Hostname:   R1
IP Address: 1.1.1.1
IPv6 Address: 1111::11
Metric:    10      IS-Extended R1.02
Metric:    10      IS (MT-IPv6) R1.02
Metric:    10      IP-Extended 1.1.1.1/32
  Prefix Attribute Flags[0]: ELC Set
Metric:    10      IP-Extended 10.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric:    10      IPv6 (MT-IPv6) 1111::11/128
Metric:    10      IPv6 (MT-IPv6) 1001::/64
R1.02-00   0x0000000C  0x724E      423      0/0/0
Metric:    0      IS-Extended R1.00
Metric:    0      IS-Extended R2.00
R2.00-00   0x00000014  0x2A52      423      0/0/0
Area Address: 49
Topology:   IPv4 (0x0) IPv6 (0x2)
NLPID:     0xCC 0x8E
Hostname:   R2
IP Address: 2.2.2.2
IPv6 Address: 2222::22
Metric:    10      IS-Extended R1.02
Metric:    10      IS-Extended R3.03
Metric:    10      IS (MT-IPv6) R1.02
Metric:    10      IS (MT-IPv6) R3.03
Metric:    10      IS (MT-IPv6) R4.04
Metric:    10      IP-Extended 2.2.2.2/32
  Prefix Attribute Flags[0]: ELC Set
Metric:    10      IP-Extended 10.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric:    10      IP-Extended 40.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric:    10      IPv6 (MT-IPv6) 2222::22/128
Metric:    10      IPv6 (MT-IPv6) 1001::/64
Metric:    10      IPv6 (MT-IPv6) 4001::/64
Metric:    10      IPv6 (MT-IPv6) 2001::/64
R3.00-00   0x00000013  0x7FCC      423      0/0/0
Area Address: 49
Topology:   IPv4 (0x0) IPv6 (0x2)
NLPID:     0xCC 0x8E
Hostname:   R3
IP Address: 3.3.3.3
IPv6 Address: 3333::33
Metric:    10      IS-Extended R4.01
Metric:    10      IS-Extended R3.03
Metric:    10      IS (MT-IPv6) R4.01
Metric:    10      IS (MT-IPv6) R3.03
Metric:    10      IP-Extended 3.3.3.3/32
Metric:    10      IP-Extended 50.1.1.0/24

```

```

Metric: 10      IP-Extended 40.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 3333::33/128
Metric: 10      IPv6 (MT-IPv6) 5001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
R3.03-00      0x0000000C  0x6D4E      423      0/0/0
Metric: 0      IS-Extended R3.00
Metric: 0      IS-Extended R2.00
R4.00-00      * 0x00000015  0x8C0D      424      0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:     R4
IP Address:   50.1.1.2
IPv6 Address: 5001::2
Metric: 10      IS-Extended R5.02
Metric: 10      IS-Extended R4.01
Metric: 10      IS (MT-IPv6) R5.02
Metric: 10      IS (MT-IPv6) R4.04
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IP-Extended 50.1.1.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 4.4.4.4/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 30.1.1.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IPv6 (MT-IPv6) 4444::44/128
Metric: 10      IPv6 (MT-IPv6) 3001::/64
Metric: 10      IPv6 (MT-IPv6) 2001::/64
Metric: 10      IPv6 (MT-IPv6) 5001::/64
R4.01-00      * 0x00000007  0x9A25      424      0/0/0
Metric: 0      IS-Extended R4.00
Metric: 0      IS-Extended R3.00
R4.04-00      * 0x0000000C  0x6751      424      0/0/0
Metric: 0      IS-Extended R4.00
Metric: 0      IS-Extended R2.00
R5.00-00      0x00000010  0xFA0F      423      0/0/0
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:     R5
IP Address:   5.5.5.5
IPv6 Address: 5555::55
Metric: 10      IS-Extended R5.02
Metric: 10      IS (MT-IPv6) R5.02
Metric: 10      IP-Extended 5.5.5.5/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IP-Extended 30.1.1.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10      IPv6 (MT-IPv6) 5555::55/128
Metric: 10      IPv6 (MT-IPv6) 3001::/64

```

```
R5.02-00          0x00000007  0xA813          423          0/0/0
Metric:  0          IS-Extended R5.00
Metric:  0          IS-Extended R4.00
```

R4#show isis database detail

Tag 1: VRF : default

IS-IS Level-1 Link State Database:

LSPID	LSP Seq Num	LSP Checksum	LSP Holdtime	ATT/P/OL
R1.00-00	0x00000015	0x9E64	417	0/0/0

```
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:    R1
IP Address:  1.1.1.1
IPv6 Address: 1111::11
Metric:  10          IS-Extended R1.02
Metric:  10          IS (MT-IPv6) R1.02
Metric:  10          IP-Extended 1.1.1.1/32
Metric:  10          IP-Extended 10.1.1.0/24
Metric:  10          IPv6 (MT-IPv6) 1111::11/128
Metric:  10          IPv6 (MT-IPv6) 1001::/64
```

R1.02-00	0x0000000C	0x724E	417	0/0/0
Metric: 0	IS-Extended R1.00			
Metric: 0	IS-Extended R2.00			

R2.00-00	0x00000014	0x2A52	417	0/0/0
----------	------------	--------	-----	-------

```
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:    R2
IP Address:  2.2.2.2
IPv6 Address: 2222::22
Metric:  10          IS-Extended R1.02
Metric:  10          IS-Extended R3.03
Metric:  10          IS (MT-IPv6) R1.02
Metric:  10          IS (MT-IPv6) R3.03
Metric:  10          IS (MT-IPv6) R4.04
Metric:  10          IP-Extended 2.2.2.2/32
Metric:  10          IP-Extended 10.1.1.0/24
Metric:  10          IP-Extended 40.1.1.0/24
Metric:  10          IPv6 (MT-IPv6) 2222::22/128
Metric:  10          IPv6 (MT-IPv6) 1001::/64
Metric:  10          IPv6 (MT-IPv6) 4001::/64
Metric:  10          IPv6 (MT-IPv6) 2001::/64
```

R3.00-00	0x00000013	0x7FCC	417	0/0/0
----------	------------	--------	-----	-------

```
Area Address: 49
Topology:     IPv4 (0x0) IPv6 (0x2)
NLPID:       0xCC 0x8E
Hostname:    R3
IP Address:  3.3.3.3
```

```

IPv6 Address: 3333::33
Metric: 10      IS-Extended R4.01
Metric: 10      IS-Extended R3.03
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IS (MT-IPv6) R3.03
Metric: 10      IP-Extended 3.3.3.3/32
Metric: 10      IP-Extended 50.1.1.0/24
Metric: 10      IP-Extended 40.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 3333::33/128
Metric: 10      IPv6 (MT-IPv6) 5001::/64
Metric: 10      IPv6 (MT-IPv6) 4001::/64
R3.03-00        0x0000000C  0x6D4E          417          0/0/0
Metric: 0      IS-Extended R3.00
Metric: 0      IS-Extended R2.00
R4.00-00        * 0x00000015  0x8C0D          418          0/0/0
Area Address: 49
Topology:      IPv4 (0x0) IPv6 (0x2)
NLPID:        0xCC 0x8E
Hostname:     R4
IP Address:   50.1.1.2
IPv6 Address: 5001::2
Metric: 10      IS-Extended R5.02
Metric: 10      IS-Extended R4.01
Metric: 10      IS (MT-IPv6) R5.02
Metric: 10      IS (MT-IPv6) R4.04
Metric: 10      IS (MT-IPv6) R4.01
Metric: 10      IP-Extended 50.1.1.0/24
Metric: 10      IP-Extended 4.4.4.4/32
Metric: 10      IP-Extended 30.1.1.0/24
Metric: 10      IPv6 (MT-IPv6) 4444::44/128
Metric: 10      IPv6 (MT-IPv6) 3001::/64
Metric: 10      IPv6 (MT-IPv6) 2001::/64
Metric: 10      IPv6 (MT-IPv6) 5001::/64
R4.01-00        * 0x00000007  0x9A25          418          0/0/0
Metric: 0      IS-Extended R4.00
Metric: 0      IS-Extended R3.00
R4.04-00        * 0x0000000C  0x6751          418          0/0/0
Metric: 0      IS-Extended R4.00
Metric: 0      IS-Extended R2.00
R5.00-00        0x00000010  0xFA0F          417          0/0/0
Area Address: 49
Topology:      IPv4 (0x0) IPv6 (0x2)
NLPID:        0xCC 0x8E
Hostname:     R5
IP Address:   5.5.5.5
IPv6 Address: 5555::55
Metric: 10      IS-Extended R5.02
Metric: 10      IS (MT-IPv6) R5.02
Metric: 10      IP-Extended 5.5.5.5/32
Metric: 10      IP-Extended 30.1.1.0/24

```

```

Metric: 10      IPv6 (MT-IPv6) 5555::55/128
Metric: 10      IPv6 (MT-IPv6) 3001::/64
R5.02-00      0x00000007  0xA813      417      0/0/0
Metric: 0      IS-Extended R5.00
Metric: 0      IS-Extended R4.00

```

R5:

R5#show clns neighbors

```

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 0
Total number of adjacencies: 1
Tag 1: VRF : default
System Id      Interface  SNPA      State  Holdtime  Type Protocol
R4             ce22/4    9819.2ccf.ede2  Up     28        L1  M-ISIS

```

R5#show clns is-neighbors detail

```

Tag 1: VRF : default
System Id      Interface  State  Type Priority  Circuit Id
R4             ce22/4    Up     L1  64        0000.0000.0005.02
L1 Adjacency ID: 1
L2 Adjacency ID: 2
Uptime: 01:13:32
Area Address(es): 49
IP Address(es): 30.1.1.1
IPv6 Address(es): fe80::9a19:2cff:fe2f:ede2
Topology: IPv4, IPv6
Level-1 Protocols Supported: IPv4, IPv6
Bidirectional Forwarding Detection is enabled
Adjacency advertisement: Advertise

```

R5#show isis topology

```

Tag 1: VRF : default
IS-IS paths to level-1 routers
System Id      Metric  Next-Hop      Interface  SNPA
R1             40     R4             ce22/4    9819.2ccf.ede2
R2             30     R4             ce22/4    9819.2ccf.ede2
R3             20     R4             ce22/4    9819.2ccf.ede2
R4             10     R4             ce22/4    9819.2ccf.ede2
R5             --

```

R5#show ipv6 isis topology

Tag 1: VRF : default

IS-IS paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
R1	30	R4 ce22/4	9819.2ccf.ede2	
R2	20	R4 ce22/4	9819.2ccf.ede2	
R3	20	R4 ce22/4	9819.2ccf.ede2	
R4	10	R4 ce22/4	9819.2ccf.ede2	
R5	--			

R5#show ip route

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

O - OSPF, IA - OSPF inter area

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

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

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

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

v - vrf leaked

* - candidate default

IP Route Table for VRF "default"

```

i L1      1.1.1.1/32 [115/50] via 30.1.1.1, ce22/4, installed 00:13:40, last update
00:13:40 ago
i L1      2.2.2.2/32 [115/40] via 30.1.1.1, ce22/4, installed 00:13:40, last update
00:13:40 ago
i L1      3.3.3.3/32 [115/30] via 30.1.1.1, ce22/4, installed 01:02:05, last update
01:02:05 ago
i L1      4.4.4.4/32 [115/20] via 30.1.1.1, ce22/4, installed 01:13:40, last update
01:13:40 ago
C         5.5.5.5/32 is directly connected, lo, installed 02:03:15, last update
02:03:15 ago
i L1      10.1.1.0/24 [115/40] via 30.1.1.1, ce22/4, installed 00:13:40, last update
00:13:40 ago
C         30.1.1.0/24 is directly connected, ce22/4, installed 02:03:15, last update
02:03:15 ago
i L1      40.1.1.0/24 [115/30] via 30.1.1.1, ce22/4, installed 01:04:55, last update
01:04:55 ago
i L1      50.1.1.0/24 [115/20] via 30.1.1.1, ce22/4, installed 01:02:05, last update
01:02:05 ago
C         127.0.0.0/8 is directly connected, lo, installed 02:20:59, last update
02:20:59 ago
  
```

Gateway of last resort is not set

R5#show ipv6 route

IPv6 Routing Table

Codes: K - kernel route, C - connected, S - static, D- DHCP, R - RIP,

O - OSPF, IA - OSPF inter area, E1 - OSPF external type 1,

E2 - OSPF external type 2, E - EVPN N1 - OSPF NSSA external type 1,

N2 - OSPF NSSA external type 2, i - IS-IS, B - BGP,

P - SRV6-POLICY,

v - vrf leaked

Timers: Uptime

IP Route Table for VRF "default"

```

C      ::1/128 via ::, lo, installed 02:21:00, last update 02:21:00 ago
i L1   1001::/64 [115/30] via fe80::9a19:2cff:febf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1   1111::11/128 [115/40] via fe80::9a19:2cff:febf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1   2001::/64 [115/20] via fe80::9a19:2cff:febf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1   2222::22/128 [115/30] via fe80::9a19:2cff:febf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
C      3001::/64 via ::, ce22/4, installed 01:05:32, last update 01:05:32 ago
i L1   3333::33/128 [115/30] via fe80::9a19:2cff:febf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1   4001::/64 [115/30] via fe80::9a19:2cff:febf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1   4444::44/128 [115/20] via fe80::9a19:2cff:febf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
i L1   5001::/64 [115/20] via fe80::9a19:2cff:febf:ede2, ce22/4, installed 00:13:41,
last update 00:13:41 ago
C      5555::55/128 via ::, lo, installed 01:06:20, last update 01:06:20 ago
C      fe80::/64 via ::, ce22/4, installed 02:03:16, last update 02:03:16 ago

```

R5#show isis spf-logs level-1-2

```

Tag 1:  VRF : default
Level-1 spf logs:
Next SPF is not scheduled yet
SPF schedule delay min 0 secs 500 msec
SPF schedule delay max 50 secs 0 msec
SPF algorithm executed 12 times
SPF algorithm last executed 00:13:45.938 ago

```

R5#show isis database verbose

```

Tag 1:  VRF : default
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
R1.00-00       0x00000015   0x9E64        373           0/0/0
Area Address:  49
Topology:      IPv4 (0x0) IPv6 (0x2)
NLPID:         0xCC 0x8E
Hostname:      R1
IP Address:    1.1.1.1
IPv6 Address:  1111::11
Metric:        10          IS-Extended R1.02
Metric:        10          IS (MT-IPv6) R1.02
Metric:        10          IP-Extended 1.1.1.1/32
Prefix Attribute Flags[0]: ELC Set
Metric:        10          IP-Extended 10.1.1.0/24
Prefix Attribute Flags[0]: ELC Set

```

```

Metric: 10 IPv6 (MT-IPv6) 1111::11/128
Metric: 10 IPv6 (MT-IPv6) 1001::/64
R1.02-00 0x0000000C 0x724E 373 0/0/0
Metric: 0 IS-Extended R1.00
Metric: 0 IS-Extended R2.00
R2.00-00 0x00000014 0x2A52 373 0/0/0
Area Address: 49
Topology: IPv4 (0x0) IPv6 (0x2)
NLPID: 0xCC 0x8E
Hostname: R2
IP Address: 2.2.2.2
IPv6 Address: 2222::22
Metric: 10 IS-Extended R1.02
Metric: 10 IS-Extended R3.03
Metric: 10 IS (MT-IPv6) R1.02
Metric: 10 IS (MT-IPv6) R3.03
Metric: 10 IS (MT-IPv6) R4.04
Metric: 10 IP-Extended 2.2.2.2/32
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 10.1.1.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IP-Extended 40.1.1.0/24
Prefix Attribute Flags[0]: ELC Set
Metric: 10 IPv6 (MT-IPv6) 2222::22/128
Metric: 10 IPv6 (MT-IPv6) 1001::/64
Metric: 10 IPv6 (MT-IPv6) 4001::/64
Metric: 10 IPv6 (MT-IPv6) 2001::/64
R3.00-00 0x00000013 0x7FCC 372 0/0/0
Area Address: 49
Topology: IPv4 (0x0) IPv6 (0x2)
NLPID: 0xCC 0x8E
Hostname: R3
IP Address: 3.3.3.3
IPv6 Address: 3333::33
Metric: 10 IS-Extended R4.01
Metric: 10 IS-Extended R3.03
Metric: 10 IS (MT-IPv6) R4.01
Metric: 10 IS (MT-IPv6) R3.03
Metric: 10 IP-Extended 3.3.3.3/32
Metric: 10 IP-Extended 50.1.1.0/24
Metric: 10 IP-Extended 40.1.1.0/24
Metric: 10 IPv6 (MT-IPv6) 3333::33/128
Metric: 10 IPv6 (MT-IPv6) 5001::/64
Metric: 10 IPv6 (MT-IPv6) 4001::/64
R3.03-00 0x0000000C 0x6D4E 372 0/0/0
Metric: 0 IS-Extended R3.00
Metric: 0 IS-Extended R2.00
R4.00-00 0x00000015 0x8C0D 373 0/0/0
Area Address: 49
Topology: IPv4 (0x0) IPv6 (0x2)

```

```

NLPID:          0xCC 0x8E
Hostname:       R4
IP Address:     50.1.1.2
IPv6 Address:   5001::2
Metric:        10          IS-Extended R5.02
Metric:        10          IS-Extended R4.01
Metric:        10          IS (MT-IPv6) R5.02
Metric:        10          IS (MT-IPv6) R4.04
Metric:        10          IS (MT-IPv6) R4.01
Metric:        10          IP-Extended 50.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric:        10          IP-Extended 4.4.4.4/32
  Prefix Attribute Flags[0]: ELC Set
Metric:        10          IP-Extended 30.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric:        10          IPv6 (MT-IPv6) 4444::44/128
Metric:        10          IPv6 (MT-IPv6) 3001::/64
Metric:        10          IPv6 (MT-IPv6) 2001::/64
Metric:        10          IPv6 (MT-IPv6) 5001::/64
R4.01-00        0x00000007  0x9A25          373          0/0/0
  Metric:        0          IS-Extended R4.00
  Metric:        0          IS-Extended R3.00
R4.04-00        0x0000000C  0x6751          373          0/0/0
  Metric:        0          IS-Extended R4.00
  Metric:        0          IS-Extended R2.00
R5.00-00        * 0x00000010  0xFA0F          373          0/0/0
Area Address:   49
Topology:       IPv4 (0x0) IPv6 (0x2)
NLPID:         0xCC 0x8E
Hostname:       R5
IP Address:     5.5.5.5
IPv6 Address:   5555::55
Metric:        10          IS-Extended R5.02
Metric:        10          IS (MT-IPv6) R5.02
Metric:        10          IP-Extended 5.5.5.5/32
  Prefix Attribute Flags[0]: ELC Set
Metric:        10          IP-Extended 30.1.1.0/24
  Prefix Attribute Flags[0]: ELC Set
Metric:        10          IPv6 (MT-IPv6) 5555::55/128
Metric:        10          IPv6 (MT-IPv6) 3001::/64
R5.02-00        * 0x00000007  0xA813          373          0/0/0
  Metric:        0          IS-Extended R5.00
  Metric:        0          IS-Extended R4.00

```

Running Configuration

```

R1#sh running-config router isis
!
router isis 1
 is-type level-1

```

```
metric-style wide
dynamic-hostname
bfd
all-interfaces
net 49.0000.0000.0001.00
!
address-family ipv6
multi-topology
level-1
exit-address-family
!
R1#
```

CLI Commands

The ISIS Multi-topology feature introduces the `multi-topology` configuration command.

multi topology

Use this command to configure the ISIS topology type.

Use `no` parameter of this command to set the topology back to single.

Command Syntax

```
multi-topology (level-1|level-1-2|level-2)
no multi-topology
```

Parameters

<code>level-1</code>	Specify to enable multi-topology for level 1
<code>level-2</code>	Specify to enable multi-topology for level 2
<code>level-1-2</code>	Specify to enable multi-topology for both the levels

Default

ISIS topology type applies to levels 1 and 2.

Command Mode

Address-family IPv6 mode.

Applicability

Introduced the `multi-topology` parameter in OcNOS version 6.5.2.

Example

The following sequence of commands is used to configure ISIS `multi-topology` type as transition for levels 1 and 2.

```
(config)#router isis 1
```

```
(config-router)#address-family ipv6 unicast
(config-router-af)#multi-topology level-1-2
```

Glossary

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

Key Terms/Acronym	Description
ISIS	Intermediate System to Intermediate System is a link-state routing protocol.
Multi Topology (MT)	In ISIS, Multi Topology (MT) is a mechanism to run a set of independent IP topologies within a single ISIS domain.
Type Length Value (TLV)	A data structure used to encode optional information in a data communications protocol: <ul style="list-style-type: none"> • Type: the kind of field that this part of the message represents • Length: the size of the value field, usually in bytes • Value: a variable-sized set of bytes that contains the data of the message
Shortest Path First (SPF)	Algorithm used by ISIS to make routing decisions based on the state of network links.
Loopback	A troubleshooting test in which a signal is transmitted from a source to a destination and then back to the source again so that the signal can be measured and evaluated.
Wide metric configuration	Allows ISIS to support larger networks by configuring high value metric in the interface.
Hello Packets	Information packets used to discover ISIS neighbors and maintain adjacencies.
Link State Packets (LSP)	Unidirectional, point-to-point, half-duplex connection used to exchange link state information.