OcNOS®

Open Compute Network Operating System for Data Center Version 6.4.2

Key Features

December 2023

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

Contents

Preface

This guide describes how to configure OcNOS.

Audience

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

Conventions

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

Table P-1: Conventions

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

Related Documentation

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

Feature Availability

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

Migration Guide

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

Support

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

Comments

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

Enhanced Security and Performance

Release 6.4.1

This section, describes the security, performance, authentication, and access control enhancements introduced in the 6.4.1 release.

- NetConf Port Access Control
- Hide the Remote AS using the neighbor local-as Command
- Support IGMP Snooping for Provider Bridge
- TCP MSS configuration for BGP neighbors
- TCP MSS configuration for LDP sessions
- Single Home VxLAN IRB with OSPF or ISIS
- Fall Back Option for RADIUS Authentication
- Modified Extended ACL Deny Rule Behavior in VTY

NetConf Port Access Control

Overview

NetConf is a software tool that provides a mechanism to configure and manage remote network devices seamlessly. It uses a simple Remote Procedure Call (RPC) mechanism to facilitate communication between a client and a server.

During the OcNOS installation, the NetConf subsystem called “netconf” is installed. It runs on the default access port 830 over SSH and port 6513 over TLS.

Typically, these default access ports are not configurable and controlled. The NetConf port access control feature enhancement ensures that the Netconf-SSH and NetConf-TLS port access can be controlled and configurable through the new CLIs introduced in the 6.4.1 release.

The following are the new CLIs introduced to support the NetConf port access control:

- feature netconf-ssh
- feature netconf-tls
- netconf-ssh port
- netconf-tls port
- show netconf server
- show running-config netconf server

The following existing CLI is updated to support the NetConf port access control

- ip access-list tcp|udp

Feature Characteristics

- This feature allows access control capabilities for the NetConf-SSH and NetConf-TLS ports.
- Enabling/disabling the port.
- Changing the default port.
- Accessing and controlling the NetConf services through Inband and Outband.
- Applying ACL rules to the NetConf port to control its access.

Benefits

This feature enables the user to control the NetConf port access and change the default port.

Configuration

To configure either NetConf-SSH port or the NetConf-TLS port, perform the following steps. After completing the steps you will be configured with a port for NetConf.

1. Disable `netconf-ssh` and `netconf-tls` feature
2. Configure port for `netconf-ssh` and `netconf-tls`
3. Enable `netconf-ssh` and `netconf-tls` feature

Topology

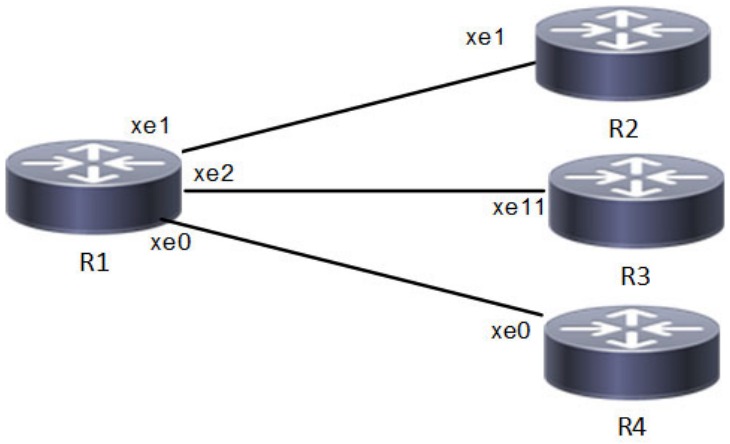

NetConf Accses Port Topology

Enable Netconf-ssh on the default and vrf management port

R1

`#configure terminal`	Enter Configuration mode.
`R1(config)#feature netconf-ssh`	Enable netconf-ssh via default port.
`R1(config)#feature netconf-ssh vrf management`	Enable netconf-ssh via vrf management port.
`R1(config)#commit`	Commit all the transactions.

Enable Netconf-tls on the default and vrf management port

R1

`#configure terminal`	Enter Configuration mode
`R1(config)#feature netconf-tls`	Enable netconf-tls via default port
`R1(config)#feature netconf-tls vrf management`	Enable netconf-tls via vrf management port
`R1(config)#commit`	Commit all the transactions

Validation

Execute the below commands to verify the NetConf port is enabled on VRF Management.

Following is the output of the NetConf server status and port.

```
#show netconf server
VRF Management
        Netconf SSH Server: Enabled
```

```
        SSH-Netconf Port : 830
        Netconf TLS Server: Enabled
        TLS-Netconf Port : 6513
VRF Default
        Netconf SSH Server: Enabled
        SSH-Netconf Port : 830
        Netconf TLS Server: Enabled
        TLS-Netconf Port : 6513
```

Following is the output of NetConf server configurations.

```
#show running-config netconf-server
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
netconf server ssh-port 2000 vrf management
netconf server tls-port 60000 vrf management
feature netconf-ssh
feature netconf-tls
netconf server ssh-port 1060
netconf server tls-port 5000
!
```

Following is the output of the NetConf server configuration in XML format.

```
#show xml running-config
<netconf-server xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-management-serve
r">
  <vrfs>
    <vrf>
      <vrf-name>default</vrf-name>
      <config>
        <vrf-name>default</vrf-name>
      </config>
      <netconf-ssh-config>
        <config>
          <feature-netconf-ssh>true</feature-netconf-ssh>
          <ssh-port>1060</ssh-port>
        </config>
      </netconf-ssh-config>
      <netconf-tls-config>
        <config>
          <feature-netconf-tls>true</feature-netconf-tls>
          <tls-port>5000</tls-port>
        </config>
      </netconf-tls-config>
    </vrf>
    <vrf>
      <vrf-name>management</vrf-name>
      <config>
        <vrf-name>management</vrf-name>
```

```
      </config>
      <netconf-ssh-config>
        <config>
          <feature-netconf-ssh>true</feature-netconf-ssh>
          <ssh-port>2000</ssh-port>
        </config>
      </netconf-ssh-config>
      <netconf-tls-config>
        <config>
          <feature-netconf-tls>true</feature-netconf-tls>
          <tls-port>60000</tls-port>
        </config>
      </netconf-tls-config>
    </vrf>
  </vrfs>
</netconf-server>
<network-instances xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-network-insta
nce">
  <network-instance>
    <instance-name>default</instance-name>
    <instance-type>vrf</instance-type>
    <config>
      <instance-name>default</instance-name>
      <instance-type>vrf</instance-type>
    </config>
    <vrf xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-vrf">
      <config>
        <vrf-name>default</vrf-name>
      </config>
    </vrf>
  </network-instance>
  <network-instance>
    <instance-name>management</instance-name>
    <instance-type>vrf</instance-type>
    <config>
      <instance-name>management</instance-name>
      <instance-type>vrf</instance-type>
    </config>
    <vrf xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-vrf">
      <config>
        <vrf-name>management</vrf-name>
      </config>
    </vrf>
  </network-instance>
</network-instances>
<interfaces xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-interface">
```

Following is the output after login to the NetConf interface (YangCLI) on R1 node via the default NetConf port:

```
root@OcNOS:~# ip netns exec zebosfib0 yangcli --server=127.1 --user=ocnos --
password=ocnos
yangcli version 2.5-5
  libssh2 version 1.8.0

  Copyright (c) 2008-2012, Andy Bierman, All Rights Reserved.
  Copyright (c) 2013-2015, Vladimir Vassilev, All Rights Reserved.
  Copyright (c) 2012-2016, OpenClovis Inc, All Rights Reserved.

  Type 'help' or 'help <command-name>' to get started
  Use the <tab> key for command and value completion
  Use the <enter> key to accept the default value in brackets

  These escape sequences are available when filling parameter values:

        ?       help
        ??      full help
        ?s      skip current parameter
        ?c      cancel current command

  These assignment statements are available when entering commands:

        $<varname> = <expr>     Local user variable assignment
        $$<varname> = <expr>    Global user variable assignment
        @<filespec> = <expr>    File assignment
val->res is NO_ERR.

yangcli: Starting NETCONF session for ocnos on 127.1

NETCONF session established for ocnos on 127.1
……..
```

Disable netconf-ssh via default and vrf management port

R1

`#configure terminal`	Enter Configuration mode
`R1(config)#no feature netconf-ssh`	Disable netconf-ssh via default port
`R1(config)#no feature netconf-ssh vrf management`	Disable netconf-ssh via vrf management
`R1(config)#commit`	Commit all the transactions

Disable netconf-tls via default port and vrf management port

R1

`#configure terminal`	Enter Configuration mode
`R1(config)#no feature netconf-tls`	Disable netconf-tls via default
`R1(config)#no feature netconf-tls vrf management`	Disable netconf-tls via vrf management port
`R1(config)#commit`	Commit all the transactions

Validation

Execute the below commands to verify the NetConf port is disabled on VRF Management.

Following is the output of the NetConf server status and port.

```
#show netconf server
VRF Management
        Netconf Server: Disabled
VRF Default
        Netconf Server: Disabled
```

Configuring NetConf Port

R1

`#configure terminal`	Enter Configuration mode
`R1(config)#no feature netconf-ssh`	Disable netconf-ssh via default port
`R1(config)#no feature netconf-ssh vrf management`	Disable netconf-ssh via vrf management port
`R1(config)#no feature netconf-tls`	Disable netconf-tls via default port
`R1(config)#no feature netconf-tls vrf management`	Disable netconf-tls via vrf management port
`R1(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R1(config)#netconf server ssh-port 1060`	Configure port for netconf-ssh default
`R1(config)#netconf server ssh-port 2000 vrf management`	Configure port for netconf-ssh vrf management
`R1(config)#netconf server tls-port 5000`	Configure port for netconf-tls default
`R1(config)#netconf server tls-port 60000 vrf management`	Configure port for netconf-tls vrf management
`R1(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R1(config)#feature netconf-ssh`	Enable netconf-ssh via default port
`R1(config)#feature netconf-ssh vrf management`	Enable netconf-ssh via vrf management port
`R1(config)#feature netconf-tls`	Enable netconf-tls via default port
`R1(config)#feature netconf-tls vrf management`	Enable netconf-tls via vrf management port
`R1(config)#commit`	Commit all the transactions

Validation

Following is the output of the NetConf server status and port.

```
#show netconf server
VRF Management
        Netconf SSH Server: Enabled
        SSH-Netconf Port : 2000
        Netconf TLS Server: Enabled
        TLS-Netconf Port : 60000
VRF Default
        Netconf SSH Server: Enabled
        SSH-Netconf Port : 1060
        Netconf TLS Server: Enabled
        TLS-Netconf Port : 5000
```

Following is the output after login to the NetConf interface (YangCLI) on R1 node via the user defined NetConf port:

```
root@OcNOS:~# ip netns exec zebosfib1 yangcli --server=127.1 --user=ocnos --
password=ocnos ncport=2000
Warning: Revision date in the future (2022-08-30), further warnings are suppressed
ietf-netconf-notifications.yang:46.4: warning(421): revision date in the future

  yangcli version 2.5-5
  libssh2 version 1.8.0

  Copyright (c) 2008-2012, Andy Bierman, All Rights Reserved.
  Copyright (c) 2013-2015, Vladimir Vassilev, All Rights Reserved.
  Copyright (c) 2012-2016, OpenClovis Inc, All Rights Reserved.

  Type 'help' or 'help <command-name>' to get started
  Use the <tab> key for command and value completion
  Use the <enter> key to accept the default value in brackets

  These escape sequences are available when filling parameter values:

        ?       help
        ??      full help
        ?s      skip current parameter
        ?c      cancel current command

  These assignment statements are available when entering commands:

        $<varname> = <expr>     Local user variable assignment
        $$<varname> = <expr>    Global user variable assignment
        @<filespec> = <expr>    File assignment
val->res is NO_ERR.

yangcli: Starting NETCONF session for ocnos on 127.1

NETCONF session established for ocnos on 127.1
.........................
Checking Server Modules...

yangcli ocnos@127.1>
```

Ping between two nodes via Yang CLI

Perform the following configurations to verify the reachability among R1, R2 and R3 routers via NetConf-SSH and NetConf-TLS port.

R1

`#configure terminal`	Enter Configuration mode
`R1(config)#no feature netconf-ssh`	Disable netconf-ssh via default
`R1(config)#no feature netconf-ssh vrf management`	Disable netconf-ssh via vrf management port
`R1(config)#no feature netconf-tls`	Disable netconf-tls via default port
`R1(config)#no feature netconf-tls vrf management`	Disable netconf-tls via vrf management port
`R1(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R1(config)#netconf server ssh-port 1060`	Configure port for netconf-ssh default
`R1(config)#netconf server ssh-port 2000 vrf management`	Configure port for netconf-ssh vrf management
`R1(config)#netconf server tls-port 5000`	Configure port for netconf-tls default
`R1(config)#netconf server tls-port 60000 vrf management`	Configure port for netconf-tls vrf management
`R1(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R1(config)#feature netconf-ssh`	Enable netconf-ssh via default port
`R1(config)#feature netconf-ssh vrf management`	Enable netconf-ssh via vrf management port
`R1(config)#feature netconf-tls`	Enable netconf-tls via default port
`R1(config)#feature netconf-tls vrf management`	Enable netconf-tls via vrf management port
`R1(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R1(config)#interface xe1`	Enter interface mode
`R1(config)#ip address 10.10.10.1/24`	Configure ipv4 address on the interface xe1.
`R1(config)#commit`	Commit all the transactions

R2

`#configure terminal`	Enter Configuration mode
`R2(config)#no feature netconf-ssh`	Disable netconf-ssh via default
`R2(config)#no feature netconf-ssh vrf management`	Disable netconf-ssh via vrf management
`R2(config)#no feature netconf-tls`	Disable netconf-tls via default
`R2(config)#no feature netconf-tls vrf management`	Disable netconf-tls via vrf management port
`R2(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R2(config)#netconf server ssh-port 1060`	Configure port for netconf-ssh default
`R2(config)#netconf server ssh-port 2000 vrf management`	Configure port for netconf-ssh vrf management
`R2(config)#netconf server tls-port 5000`	Configure port for netconf-tls default
`R2(config)#netconf server tls-port 60000 vrf management`	Configure port for netconf-tls vrf management
`R2(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R2(config)#feature netconf-ssh`	Enable netconf-ssh via default port
`R2(config)#feature netconf-ssh vrf management`	Enable netconf-ssh via vrf management port
`R2(config)#feature netconf-tls`	Enable netconf-tls via default port
`R2(config)#feature netconf-tls vrf management`	Enable netconf-tls via vrf management port
`R2(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R2(config)#interface xe1`	Enter interface mode
`R2(config)#ip address 10.10.10.2/24`	Configure ipv4 address on the interface xe1.
`R2(config)#commit`	Commit all the transactions

Validation

Following is the output of the configured NetConf port.

```
#show netconf server
VRF Management
        Netconf SSH Server: Enabled
        SSH-Netconf Port : 2000
        Netconf TLS Server: Enabled
        TLS-Netconf Port : 60000
VRF Default
        Netconf SSH Server: Enabled
        SSH-Netconf Port : 1060
        Netconf TLS Server: Enabled
        TLS-Netconf Port : 5000
```

```
OcNOS#show running-config interface xe1
!
interface xe1
 ip address 10.10.10.1/24
!
OcNOS#ping 10.10.10.2
Press CTRL+C to exit
PING 10.10.10.2 (10.10.10.2) 56(84) bytes of data.
64 bytes from 10.10.10.2: icmp_seq=1 ttl=64 time=0.567 ms
64 bytes from 10.10.10.2: icmp_seq=2 ttl=64 time=0.258 ms
64 bytes from 10.10.10.2: icmp_seq=3 ttl=64 time=0.241 ms

--- 10.10.10.2 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 80ms
rtt min/avg/max/mdev = 0.241/0.355/0.567/0.150 ms
```

Following is the output after login to the NetConf interface (YangCLI) on R2 node through the user defined NetConf port:

```
root@OcNOS:~# ip netns exec zebosfib0 yangcli --server=10.10.10.2 --user=ocnos --
password=ocnos ncport=1060
Warning: Revision date in the future (2022-08-30), further warnings are suppressed
ietf-netconf-notifications.yang:46.4: warning(421): revision date in the future

  yangcli version 2.5-5
  libssh2 version 1.8.0

  Copyright (c) 2008-2012, Andy Bierman, All Rights Reserved.
  Copyright (c) 2013-2015, Vladimir Vassilev, All Rights Reserved.
  Copyright (c) 2012-2016, OpenClovis Inc, All Rights Reserved.

  Type 'help' or 'help <command-name>' to get started
  Use the <tab> key for command and value completion
  Use the <enter> key to accept the default value in brackets

  These escape sequences are available when filling parameter values:

        ?       help
        ??      full help
        ?s      skip current parameter
        ?c      cancel current command

  These assignment statements are available when entering commands:

        $<varname> = <expr>     Local user variable assignment
        $$<varname> = <expr>    Global user variable assignment
        @<filespec> = <expr>    File assignment
val->res is NO_ERR.
```

```
yangcli: Starting NETCONF session for ocnos on 10.10.10.2

NETCONF session established for ocnos on 10.10.10.2
.........................
Checking Server Modules...

yangcli ocnos@10.10.10.2>
```

ACL Rule with IPv4 Configuration

Perform the following configurations to apply an ACL rule to allow or deny traffic from R1 to other nodes via NetConf port.

R1

`#configure terminal`	Enter Configuration mode
`R1(config)#no feature netconf-ssh`	Disable netconf-ssh via default
`R1(config)#no feature netconf-ssh vrf management`	Disable netconf-ssh via vrf management port
`R1(config)#no feature netconf-tls`	Disable netconf-tls via default port
`R1(config)#no feature netconf-tls vrf management`	Disable netconf-tls via vrf management port
`R1(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R1(config)#netconf server ssh-port 1060`	Configure port for netconf-ssh default
`R1(config)#netconf server ssh-port 2000 vrf management`	Configure port for netconf-ssh vrf management
`R1(config)#netconf server tls-port 5000`	Configure port for netconf-tls default
`R1(config)#netconf server tls-port 60000 vrf management`	Configure port for netconf-tls vrf management
`R1(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R1(config)#feature netconf-ssh`	Enable netconf-ssh via default port
`R1(config)#feature netconf-ssh vrf management`	Enable netconf-ssh via vrf management port
`R1(config)#feature netconf-tls`	Enable netconf-tls via default port
`R1(config)#feature netconf-tls vrf management`	Enable netconf-tls via vrf management port
`R1(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R1(config)#interface xe1`	Enter interface mode

R1(config)#ip address 10.10.10.1/24	Configure ipv4 address on the interface xe1.
R1(config)#commit	Commit all the transactions
#configure terminal	Enter Configuration mode
R1(config)#interface xe2	Enter interface mode
R1(config)#ip address 20.20.20.1/24	Configure ipv4 address on the interface xe2.
R1(config)#commit	Commit all the transactions
#configure terminal	Enter Configuration mode
R1(config)#ip access-list ACL1	Create ip access list
R1(config)#permit any host 10.1.1.1 any	Create an acl rule to permit
R1(config)#deny any host 20.1.1.1 any	Create an acl rule to deny
R1(config)#commit	Commit all the transactions

R2

Perform the following configurations to apply an ACL rule to allow or deny traffic from R2 to other nodes via NetConf port

#configure terminal	Enter Configuration mode
R2(config)#no feature netconf-ssh	Disable netconf-ssh via default
R2(config)#no feature netconf-ssh vrf management	Disable netconf-ssh via vrf management
R2(config)#no feature netconf-tls	Disable netconf-tls via default
R2(config)#no feature netconf-tls vrf management	Disable netconf-tls via vrf management port
R2(config)#commit	Commit all the transactions
#configure terminal	Enter Configuration mode
R2(config)#netconf server ssh-port 1060	Configure port for netconf-ssh default
R2(config)#netconf server ssh-port 2000 vrf management	Configure port for netconf-ssh vrf management
R2(config)#netconf server tls-port 5000	Configure port for netconf-tls default
R2(config)#netconf server tls-port 60000 vrf management	Configure port for netconf-tls vrf management
R2(config)#commit	Commit all the transactions
#configure terminal	Enter Configuration mode
R2(config)#feature netconf-ssh	Enable netconf-ssh via default port
R2(config)#feature netconf-ssh vrf management	Enable netconf-ssh via vrf management port
R2(config)#feature netconf-tls	Enable netconf-tls via default port
R2(config)#feature netconf-tls vrf management	Enable netconf-tls via vrf management port

`R2(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R2(config)#interface xe1`	Enter interface mode
`R2(config)#ip address 10.10.10.2/24`	Configure ipv4 address on the interface xe1.
`R2(config)#commit`	Commit all the transactions

R3

Perform the following configurations to apply an ACL rule to allow or deny traffic from R3 to other nodes via NetConf port.

`#configure terminal`	Enter Configuration mode
`R3(config)#no feature netconf-ssh`	Disable netconf-ssh via default
`R3(config)#no feature netconf-ssh vrf management`	Disable netconf-ssh via vrf management port
`R3(config)#no feature netconf-tls`	Disable netconf-tls via default port
`R3(config)#no feature netconf-tls vrf management`	Disable netconf-tls via vrf management port
`R3(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R3(config)#netconf server ssh-port 1060`	Configure port for netconf-ssh default
`R3(config)#netconf server ssh-port 2000 vrf management`	Configure port for netconf-ssh vrf management
`R3(config)#netconf server tls-port 5000`	Configure port for netconf-tls default
`R3(config)#netconf server tls-port 60000 vrf management`	Configure port for netconf-tls vrf management
`R3(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R3(config)#feature netconf-ssh`	Enable netconf-ssh via default port
`R3(config)#feature netconf-ssh vrf management`	Enable netconf-ssh via vrf management port
`R3(config)#feature netconf-tls`	Enable netconf-tls via default port
`R3(config)#feature netconf-tls vrf management`	Enable netconf-tls via vrf management port
`R3(config)#commit`	Commit all the transactions
`#configure terminal`	Enter Configuration mode
`R3(config)#interface xe11`	Enter interface mode
`R3(config)#ip address 20.20.20.2/24`	Configure ipv4 address on the interface xe11.
`R3(config)#commit`	Commit all the transactions

Validation

Following is the output to verify the user defined NetConf port.

```
R1#show running-config netconf-server
!
feature netconf-ssh vrf management
feature netconf-tls vrf management
netconf server ssh-port 2000 vrf management
netconf server tls-port 60000 vrf management
feature netconf-ssh
feature netconf-tls
netconf server ssh-port 1060
netconf server tls-port 5000
!

R1#show netconf server
VRF Management
        Netconf SSH Server: Enabled
        SSH-Netconf Port : 2000
        Netconf TLS Server: Enabled
        TLS-Netconf Port : 60000
VRF Default
        Netconf SSH Server: Enabled
        SSH-Netconf Port : 1060
        Netconf TLS Server: Enabled
        TLS-Netconf Port : 5000
```

Following is the output of the show running-config in XML format.

```
R1#show xml running-config
<netconf-server xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-management-serve
r">
  <vrfs>
    <vrf>
      <vrf-name>default</vrf-name>
      <config>
        <vrf-name>default</vrf-name>
      </config>
      <netconf-ssh-config>
        <config>
          <feature-netconf-ssh>true</feature-netconf-ssh>
          <ssh-port>1060</ssh-port>
        </config>
      </netconf-ssh-config>
      <netconf-tls-config>
        <config>
          <feature-netconf-tls>true</feature-netconf-tls>
          <tls-port>5000</tls-port>
        </config>
      </netconf-tls-config>
    </vrf>
    <vrf>
      <vrf-name>management</vrf-name>
```

```
      <config>
        <vrf-name>management</vrf-name>
      </config>
      <netconf-ssh-config>
        <config>
          <feature-netconf-ssh>true</feature-netconf-ssh>
          <ssh-port>2000</ssh-port>
        </config>
      </netconf-ssh-config>
      <netconf-tls-config>
        <config>
          <feature-netconf-tls>true</feature-netconf-tls>
          <tls-port>60000</tls-port>
        </config>
      </netconf-tls-config>
    </vrf>
  </vrfs>
</netconf-server>
<network-instances xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-network-insta
nce">
  <network-instance>
    <instance-name>default</instance-name>
    <instance-type>vrf</instance-type>
    <config>
      <instance-name>default</instance-name>
      <instance-type>vrf</instance-type>
    </config>
    <vrf xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-vrf">
      <config>
        <vrf-name>default</vrf-name>
      </config>
    </vrf>
  </network-instance>
  <network-instance>
    <instance-name>management</instance-name>
    <instance-type>vrf</instance-type>
    <config>
      <instance-name>management</instance-name>
      <instance-type>vrf</instance-type>
    </config>
    <vrf xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-vrf">
      <config>
        <vrf-name>management</vrf-name>
      </config>
    </vrf>
  </network-instance>
</network-instances>
<interfaces xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-interface">
```

Implementation Examples

The below examples are based on the topology given in Topology section.

Accessing R1 from R2 with default port

Below is an example to access R1 from R2 with default port.

From OcNOS CLI:

```
feature netconf-ssh
feature netconf-ssh vrf management
feature netconf-tls
feature netconf-tls vrf management
```

From Yang CLI:

```
root@OcNOS:~# ip netns exec zebosfib0 yangcli --server=127.1 --user=ocnos --
password=ocnos
```

Accessing R1 from R2 with user defined port

Below is an example to access R1 from R2 via user defined port.

From OcNOS CLI:

```
netconf server ssh-port 1060
netconf server ssh-port 2000 vrf management
netconf server tls-port 5000
netconf server tls-port 60000 vrf management
```

From Yang CLI:

```
root@OcNOS:~#ip netns exec zebosfib1 yangcli --server=10.10.10.1 --user=ocnos --
password=ocnos ncport=2000
```

Applying ACL rule to permit or deny any Node

Below is an example to permit any traffic originating from IP address 10.1.1.1. and deny any traffic originating from 20.1.1.1.

From OcNOS CLI:

```
ip access-list ACL1
permit any host 10.1.1.1  any
deny any host   20.1.1.1 any
Permitting R2 and denying R3
```

From Yang CLI:

```
root@OcNOS:~# ip netns exec zebosfib1 yangcli --server=10.10.10.2 --user=ocnos --
password=ocnos ncport=2000
```

New CLI Commands

feature netconf-ssh

Use this command to enable or disable the netconf-ssh feature specific to the management VRF. When netconf feature-ssh is enabled, it allows the logins through the default netconf-ssh port or through default ssh port if feature SSH is also enabled.

Command Syntax

```
feature netconf-ssh (vrf management|)
no feature netconf-ssh (vrf management|)
```

Parameters

`vrf management` Specifies the management Virtual Routing and Forwarding

Default

Disabled by default.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The following example shows you how to enable NetConf SSH on either the VRF management port or the default port. The no parameter disables the same.

```
(config)#feature netconf-ssh vrf management
(config)#feature netconf-ssh
(config)#no feature netconf-ssh vrf management
(config)#no feature netconf-ssh
#
```

feature netconf-tls

Use this command to enable or disable the NetConf TLS feature specific to a VRF. When netconf feature-ssh is enabled, it allows the logins through the default netconf-tls port and allows login through a default TLS port when the TLS feature is also enabled.

Command Syntax

```
feature netconf-tls (vrf management|)
no feature netconf-tls (vrf management|)
```

Parameters

`vrf management` Specifies management Virtual Routing and Forwarding.

Default

Disabled by default.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The following example shows how to execute the CLI:

```
(config)#feature netconf-tls vrf management
(config)#feature netconf-tls
(config)#no feature netconf-tls vrf management
(config)#no feature netconf-tls
```

If either NetConf SSH or NetConf TLS are disabled one after the other, the following error message will be displayed, `% Disabling this will stop the netconf service that is running in management vrf”` as shown below.

Management VRF Configuration

```
OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#no feature netconf-ssh vrf management
(config)#commit
(config)#no feature netconf-tls vrf management
(config)#commit
% Disabling this will stop the netconf service that is running in management vrf.
```

Default VRF Configuration

```
OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
(config)#no feature netconf-ssh vrf management
(config)#commit
(config)#no feature netconf-tls vrf management
(config)#commit
% Disabling this will stop the netconf service that is running in defualt vrf.
```

netconf-ssh port

Use this command to either configure or unconfigure the custom NetConf SSH port.

Command Syntax

```
netconf-server ssh-port <1024-65535> (vrf management|)
```

```
no netconf-server ssh-port (vrf management|)
```

Parameters

<1024-65535>	Port range values
Default	By default, the netconf-ssh port value is 830.
vrf	Specifies the management Virtual Routing and Forwarding name

Command Mode

Config mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The following example shows how to execute the CLI:

```
(config)#netconf server ssh-port ?
  <1024-65535>  port
(config)#netconf server ssh-port 1024 vrf management
(config)#netconf server ssh-port 2000
(config)#no netconf server ssh-port
(config)#no netconf server ssh-port vrf management
```

netconf-tls port

Use this command to either configure or unconfigure the indicated NetConf TLS port.

Command Syntax

```
netconf-server tls-port <1024-65535> (vrf management|)
no netconf-server tls-port (vrf management|)
```

Parameters

<1024-65535>	Port range values
Default	By default, the netconf-tls port value is 6513.
vrf	Specifies the management Virtual Routing and Forwarding name

Command Mode

Config mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
(config)#netconf server tls-port ?
  <1024-65535>  port
(config)#netconf server tls-port 5000 vrf management
(config)#netconf server tls-port 3000
(config)#no netconf server tls-port vrf management
```

```
(config)#no netconf server tls-port
```

show netconf server

Use this command to display netconf server status.

Command Syntax

```
show netconf server
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

The following example shows the output of the CLI:

```
OcNOS#show netconf server
VRF MANAGEMENT
Netconf Server: Enabled
SSH-Netconf Port : 1000
TLS-Netconf Port : 7000
VRF DEFAULT
Netconf Server: Enabled
SSH-Netconf Port : 4500
TLS-Netconf Port : 3000
```

show running-config netconf server

Use this command to display the NetConf server settings that appear in the running configuration.

Command Syntax

```
show running-config netconf-server
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The following example shows the output of the CLI:

```
OcNOS#show running-config netconf-server
feature netconf vrf management
netconf server ssh-port 1000 vrf management
netconf server tls-port 7000 vrf management
feature netconf
netconf server ssh-port 4500
netconf server tls-port 3000
!
```

Revised CLI Commands

The existing `ip access-list tcp|udp` CLI is updated with the following two options to support the Access List (ACL) rules on the NetConf port. The ACL defines a set of rules to control network traffic and reduce network attacks.

`netconf-ssh`	Secure Shell Network Configuration
`netconf-tls`	Transport Layer Security Network Configuration

ip access-list tcp|udp

Use this command to define a named access control list (ACL) that determines whether to accept or drop an incoming TCP or UDP IP packet based on the specified match criteria. This command filters packets based on source and destination IP address along with the TCP or UDP protocol and port.

Use the `no` form of this command to remove an ACL specification.

Note: Configuring same filter again with change of sequence number or change of action will result in update of sequence number or filter action.

Note: TCP flags options and range options like neq, gt, lt and range are not supported by hardware in egress direction.

Note: Both Ack and established flag in tcp have same functionality in hardware.

Command Syntax

```
(<1-268435453>|) (deny|permit) tcp (A.B.C.D/M|A.B.C.D A.B.C.D|host A.B.C.D|any)
  ((eq|gt|lt|neq) (<0-65535>|bgp|chargen|cmd|daytime|discard|domain|drip|echo
  |exec|finger|ftp |ftp-data|gopher|hostname|ident|irc|klogin|kshell|login
  |lpd|nntp|pim-auto- rp|pop2|pop3|smtp|ssh|sunrpc|tacacs|talk|telnet|time|
  uucp|whois|www)| range <0-65535> <0-65535>|) (A.B.C.D/M|A.B.C.D A.B.C.D|host
  A.B.C.D|any) ((eq|gt|lt|neq) (<0-65535>|bgp|chargen|cmd|daytime|discard|domain|
  drip|echo|exec|finger|ftp|ftp-data|gopher|hostname|ident|irc|klogin|kshell|login
  |lpd|nntp|pim-auto- rp|pop2|pop3|smtp|ssh|sunrpc|tacacs|talk|telnet
  |time|uucp|whois|www|netconf-ssh|netconf-tls) | range <0-65535> <0-65535>|)
  ((dscp (<0-63>| af11| af12| af13| af21| af22| af23| af31| af32| af33| af41| af42|
  af43| cs1| cs2| cs3| cs4| cs5| cs6| cs7| default| ef)) |(precedence (<0-7>|
  critical| flash | flashoverride| immediate| internet| network| priority|
  routine)) |) ({ack|established|fin|psh|rst|syn|urg}|) vlan <1-4094>|)(inner-vlan
  <1-4094>|)

(<1-268435453>|) (deny|permit) udp (A.B.C.D/M|A.B.C.D A.B.C.D|host A.B.C.D|any)
  ((eq|gt|lt|neq) (<0-65535>|biff|bootpc|bootps|discard|dnsix|domain|
  echo|isakmp|mobile-ip |nameserver | netbios-dgm | netbios-ns| netbios-ss|non500-
```

```
  isakmp|ntp|pim-auto-rp|rip|snmp|snmptrap|sunrpc|syslog|tacacs|talk|tftp
  |time|who|xdmcp) | range <0-65535> <0-65535>|) (A.B.C.D/M|A.B.C.D A.B.C.D|host
  A.B.C.D|any) ((eq|gt |lt|neq)(<0-65535> |biff |bootpc |bootps| discard| dnsix|
  domain| echo| isakmp|mobile-ip|nameserver|netbios-dgm|netbios-ns|netbios-
  ss|non500-isakmp |ntp|pim-auto- rp| rip| snmp| snmptrap| sunrpc| syslog| tacacs|
  talk| tftp| time| who| xdmcp) | range <0-65535> <0-65535>|) ((dscp (<0-63>| af11|
  af12| af13| af21| af22| af23| af31| af32| af33| af41| af42| af43| cs1| cs2| cs3|
  cs4| cs5| cs6| cs7| default| ef)) | (precedence (<0-7>| critical| flash |
  flashoverride| immediate| internet| network| priority| routine))|) (vlan <1-
  4094>|)(inner-vlan <1-4094>|)
no (<1-268435453>|) (deny|permit) tcp (A.B.C.D/M|A.B.C.D A.B.C.D|host
  A.B.C.D|any)((eq|gt|lt|neq) (<0-65535>| bgp| chargen| cmd| daytime| discard|
  domain| drip| echo|exec|finger|ftp |ftp-data |gopher |hostname| ident| irc|
  klogin| kshell|login|lpd|nntp|pim-auto-rp |pop2 |pop3 |smtp| ssh| sunrpc| tacacs
  |talk|telnet|time|uucp|whois|www|netconf-ssh|netconf-tls) | range <0-65535> <0-
  65535>|) (A.B.C.D/M|A.B.C.D A.B.C.D|host A.B.C.D|any)((eq|gt|lt|neq) (<0-65535>
  |bgp |chargen |cmd |daytime|discard|domain|drip|echo|exec|finger|ftp|ftp-data|
  gopher| hostname| ident| irc| klogin| kshell| login| lpd| nntp| pim-auto-rp |
  pop2| pop3| smtp |ssh |sunrpc|tacacs|talk|telnet|time|uucp|whois|www) | range <0-
  65535> <0-65535>|) ((dscp (<0-63>| af11| af12| af13| af21| af22| af23| af31|
  af32| af33| af41| af42| af43| cs1| cs2| cs3| cs4| cs5| cs6| cs7| default| ef)) |
  (precedence (<0-7>| critical| flash | flashoverride| immediate| internet|
  network| priority| routine)) |) ({ack|established|fin|psh|rst|syn|urg}|)(vlan <1-
  4094>|)(inner-vlan <1-4094>|)
no (<1-268435453>|)(deny|permit) udp (A.B.C.D/M|A.B.C.D A.B.C.D|host A.B.C.D|any)
  ((eq|gt|lt|neq) (<0-65535> |biff| bootpc| bootps| discard| dnsix|
  domain|echo|isakmp|mobile-ip|nameserver|netbios-dgm|netbios-ns|netbios-
  ss|non500-isakmp|ntp|pim-auto-rp|rip|snmp|snmptrap|sunrpc|syslog|tacacs|talk|
  tftp|time|who|xdmcp) | range <0-65535> <0-65535>|)(A.B.C.D/M|A.B.C.D A.B.C.D|host
  A.B.C.D| any) ((eq|gt|lt|neq) (<0-65535> |biff| bootpc| bootps| discard| dnsix|
  domain|echo| isakmp|mobile- ip|nameserver|netbios-dgm|netbios-ns|netbios-
  ss|non500-isakmp| ntp|pim-auto-rp|rip|snmp|snmptrap|sunrpc|syslog|
  tacacs|talk|tftp|time|who|xdmcp) | range <0-65535> <0-65535>|) ((dscp (<0-63>|
  af11| af12| af13| af21| af22| af23| af31| af32| af33| af41| af42| af43| cs1| cs2|
  cs3| cs4| cs5| cs6| cs7| default| ef)) | (precedence (<0-7>| critical| flash |
  flashoverride| immediate| internet| network| priority| routine)) |)(vlan <1-
  4094>|)(inner-vlan <1-4094>|)
```

Parameters

`<1-268435453>`	IPv4 ACL sequence number.
`deny`	Drop the packet.
`permit`	Accept the packet.
`tcp`	Transmission Control Protocol.
`udp`	User Datagram Protocol.
`A.B.C.D/M`	Source or destination IP prefix and length.
`A.B.C.D A.B.C.D`	Source or destination IP address and mask.
`host A.B.C.D`	Source or destination host IP address.
`any`	Any source or destination IP address.

`eq`	Source or destination port equal to.
`gt`	Source or destination port greater than.
`lt`	Source or destination port less than.
`neq`	Source or destination port not equal to.
`<0-65535>`	Source or destination port number.
`range`	Range of source or destination port numbers:
`<0``-65535>`	Lowest value in the range.
`<0-`**`65535>`**	Highest value in the range.
`bgp`	Border Gateway Protocol.
`chargen`	Character generator.
`cmd`	Remote commands.
`daytime`	Daytime.
`discard`	Discard.
`domain`	Domain Name Service.
`drip`	Dynamic Routing Information Protocol.
`echo`	Echo.
`exec`	EXEC.
`finger`	Finger.
`ftp`	File Transfer Protocol.
`ftp-data`	FTP data connections.
`gopher`	Gopher.
`hostname`	NIC hostname server.
`ident`	Ident Protocol.
`irc`	Internet Relay Chat.
`klogin`	Kerberos login.
`kshell`	Kerberos shell.
`login`	Login.
`lpd`	Printer service.
`nntp`	Network News Transport Protocol.
`pim-auto-rp`	PIM Auto-RP.
`pop2`	Post Office Protocol v2.
`pop3`	Post Office Protocol v3.
`smtp`	Simple Mail Transport Protocol.
`ssh`	Secure Shell.
`sunrpc`	Sun Remote Procedure Call.
`tacacs`	TAC Access Control System.
`talk`	Talk.
`telnet`	Telnet.
`time`	Time.

uucp	UNIX-to-UNIX Copy Program.
whois	WHOIS/NICNAME
www	World Wide Web.
netconf-ssh	**Secure Shell Network Configuration**
netconf-tls	**Transport Layer Security Network Configuration**
nntp	Range of source or destination port numbers:
dscp	Match packets with given DSCP value.
<0-63>	Enter DSCP value between 0-63.
af11	AF11 DSCP (001010) decimal value 10.
af12	AF12 DSCP (001100) decimal value 12.
af13	AF13 DSCP (001110) decimal value 14.
af21	AF21 DSCP (010010) decimal value 18.
af22	AF22 DSCP (010100) decimal value 20.
af23	AF23 DSCP (010110) decimal value 22.
af31	AF31 DSCP (011010) decimal value 26.
af32	AF32 DSCP (011100) decimal value 28.
af33	AF33 DSCP (011110) decimal value 30.
af41	AF41 DSCP (100010) decimal value 34.
af42	AF42 DSCP (100100) decimal value 36.
af43	AF43 DSCP (100110) decimal value 38.
cs1	CS1 (precedence 1) DSCP (001000) decimal value 8.
cs2	CS2 (precedence 2) DSCP (010000) decimal value 16.
cs3	CS3 (precedence 3) DSCP (011000) decimal value 24.
cs4	CS4 (precedence 4) DSCP (100000) decimal value 32.
cs5	CS5 (precedence 5) DSCP (101000) decimal value 40.
cs6	CS6 (precedence 6) DSCP (110000) decimal value 48.
cs7	CS7 (precedence 7) DSCP (111000) decimal value 56.
default	Default DSCP (000000) decimal value 0.
ef	EF DSCP (101110) decimal value 46.
precedence	Match packets with given precedence value.
<0-7>	Precedence.
critical	Match packets with critical precedence (5).
flash	Match packets with flash precedence (3).
flashoverride	Match packets with flash override precedence (4).
immediate	Match packets with immediate precedence (2).
internet	Match packets with internetwork control precedence (6).
network	Match packets with network control precedence (7).
priority	Match packets with priority precedence (1).
routine	Match packets with routine precedence (0).

`ack`	Match on the Acknowledgment (ack) bit.
`established`	Matches only packets that belong to an established TCP connection.
`fin`	Match on the Finish (fin) bit.
`psh`	Match on the Push (psh) bit.
`rst`	Match on the Reset (rst) bit.
`syn`	Match on the Synchronize (syn) bit.
`urg`	Match on the Urgent (urg) bit.
`biff`	Biff.
`bootpc`	Bootstrap Protocol (BOOTP) client.
`bootps`	Bootstrap Protocol (BOOTP) server.
`discard`	Discard.
`dnsix`	DNSIX security protocol auditing.
`domain`	Domain Name Service.
`echo`	Echo.
`isakmp`	Internet Security Association and Key Management Protocol.
`mobile-ip`	Mobile IP registration.
`nameserver`	IEN116 name service.
`netbios-dgm`	Net BIOS datagram service.
`netbios-ns`	Net BIOS name service.
`netbios-ss`	Net BIOS session service.
`non500-isakmp`	Non500-Internet Security Association and Key Management Protocol.
`ntp`	Network Time Protocol.
`pim-auto-rp`	PIM Auto-RP.
`rip`	Routing Information Protocol.
`snmp`	Simple Network Management Protocol.
`snmptrap`	SNMP Traps.
`sunrpc`	Sun Remote Procedure Call.
`syslogS`	ystem Logger.
`tacacs`	TAC Access Control System.
`talk`	Talk.
`tftp`	Trivial File Transfer Protocol.
`time`	Time.
`who`	Who service.
`xdmcp`	X Display Manager Control Protocol.
`vlan`	Match packets with given vlan value.
`<1-4094>`	VLAN identifier.
`inner-vlan`	Match packets with given inner vlan value.
`<1-4094>`	VLAN identifier.

Default

No default value is specified.

Command Mode

IP access-list mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

The following is an example to execute the CLI:

```
#configure terminal
(config)#ip access-list ip-acl-02
(config-ip-acl)#deny udp any any eq tftp
(config-ip-acl)#deny tcp any any eq ssh
(config-ip-acl)#end.
```

Abbreviations

Acronym	Description
ACL	Access control list
RPC	Remote Procedure Call
SSH	Secure Shell
TLS	Transport Layer Security

Hide the Remote AS using the neighbor local-as Command

Overview

In a network, an Autonomous System (AS) is available to define a set of IP routing prefixes that are under a common administration policy control. These defined routing policies are used by other connected routers o n the Internet. When an AS is configured in Border Gateway Protocol (BGP), it is used to share routing information to connected peers. The `neighbor local-as` CLI command configures the AS number to be used with External Border Gateway Protocol (EBGP) peers. By default, the configured AS number is included in the AS-PATH message that is exchanged between the peers.

When a BGP router, configured in one network, connects to another router on the network, it will automatically share routing information with the AS number of both the local and remote routers in the AS-PATH message with other connected, external peers. For example, if a router ISP1-R, accesses services from another router, ISP2-R, ISP1-R router will share routing information with local and remote AS numbers in the AS-PATH message when services are merged. This allows the external peers to learn the AS numbers of remote routers not connected to it (in this case, the AS number of ISP2-R). It is not desirable to disclose the AS number of remote routers to external peers.

To avoid advertising the remote peer's AS number, OcNOS provides an option in the `neighbor local-as` CLI to not include (`no-prepend`) the remote AS number and replace (`replace-as`) it with alternate AS number. Configuring an alternate AS in the BGP neighbor system, provides the ability to hide the AS number of the remote router that actually shares the services. Thus, the AS number of the BGP router that is actually providing services is unknown to the external peer.

Hence, the existing `neighbor local-as` CLI command has been modified in this release.

Feature Characteristics

The `neighbor local-as` CLI is enhanced to hide and replace the AS number of the remote routers not connected to external peer. Two new options '`no-prepend`' and '`replace-as`' have been added. These options replace the AS number with an alternate AS number in the AS_PATH and BGP OPEN message. Hence, the AS of the remote router is unknown to the respective neighbor peer.

Benefits

The actual Autonomous System number is never shared to the external network.

Configuration

The following configuration assumes the router R1 and R2 is assigned with AS300 and AS100 respectively.

Topology

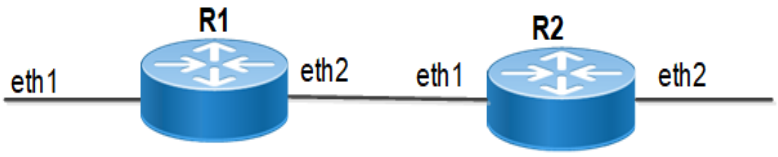

Disparate Autonomous System Number

R1

Perform the following configuration on R1 router.

`#configure terminal`	Enter configure mode.
`R1(config)#router bgp 300`	Start the BGP process with the Autonomous System number 300
`R1(config-router)#neighbor 10.10.10.2 remote-as 200`	Establish BGP session with neighbor that has AS number 200
`R1(config-router)#address-family ipv4 unicast`	Enter address-family ipv4 unicast mode
`R1(config-router-af)#neighbor 10.10.10.2 activate`	Enable the neighbor 10.10.10.2 router to exchange address family routes
`R1(config-router-af)#redistribute connected`	Redistribute information from connected routes
`R1(config-router-af)#exit-address-family`	Exit address-family IPv4 unicast mode
`R1(config-router)#commit`	Commit the configurations

R2

Perform the following configuration on R2 router.

`#configure terminal`	Enter configure mode
`R2(config)#router bgp 100`	Start the BGP process with the Autonomous System number 100
`R2(config-router)#neighbor 10.10.10.1 remote-as 300`	Establish BGP session with neighbor 10.10.10.1 that has AS number 300
`R2(config-router)#neighbor 10.10.10.1 local-as 200 no-prepend replace-as`	Replace the AS number 300 with AS number 200 that should be used with the neighbor 10.10.10.1
`R2(config-router)#address-family ipv4 unicast`	Enable the neighboring router to exchange address family routes
`R2(config-router-af)#neighbor 10.10.10.2 activate`	Enable the neighbor 10.10.10.2 router to exchange address family routes
`R2(config-router-af)#redistribute connected`	Redistribute information from the connected routes
`R2(config-router-af)#exit-address-family`	Exit address-family ipv4 unicast mode
`R2(config-router)#commit`	Commit the configurations

Validation

Check the AS number 300 running on R1. It has established a BGP connection with 10.10.10.2 router that has AS number of 200.

R1

```
OcNOS#show running-config bgp
!
router bgp 300
 neighbor 10.10.10.2 remote-as 200
 !
 address-family ipv4 unicast
 redistribute connected
 redistribute static
 neighbor 10.10.10.2 activate
 exit-address-family
!
OcNOS#
OcNOS#show ip bgp summary
BGP router identifier 10.10.10.1, local AS number 300
BGP table version is 4
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd
10.10.10.2               4   200  185        181       3      0      0  00:00:28       2

Total number of neighbors 1

Total number of Established sessions 1
OcNOS#

OcNOS#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            10.10.10.0/24 is directly connected, ce1, 1d14h18m
B            30.30.30.0/24 [20/0] via 10.10.10.2, ce1, 00:00:18
C            40.40.40.0/24 is directly connected, xe33, 1d13h40m
C            127.0.0.0/8 is directly connected, lo, 1d14h23m
Gateway of last resort is not set
OcNOS#
```

Check if the AS number 100 for R2 has been replaced with AS number 200 before sharing the information with R1.

R2

```
OcNOS#show running-config bgp
!
router bgp 100
 neighbor 10.10.10.1 remote-as 300
 neighbor 10.10.10.1 local-as 200
 !
 address-family ipv4 unicast
 redistribute connected
 redistribute static
 neighbor 10.10.10.1 activate
 exit-address-family
!
OcNOS#
OcNOS#show ip bgp summary
BGP router identifier 10.10.10.2, local AS number 100
BGP table version is 2
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Down   State/
PfxRcd
10.10.10.1              4   300  180       186      2      0      0       00:00:39       2

Total number of neighbors 1

Total number of Established sessions 1
```

Check if the AS number for R2 is changed to 100 and R1 shares AS 100 in the AS-PATH message.

R1

```
OcNOS#
OcNOS#
OcNOS#show ip bgp
BGP table version is 4, local router ID is 10.10.10.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric    LocPrf   Weight Path
*> 10.10.10.0/24    0.0.0.0             0         100      32768  ?
*                   10.10.10.2          0         100       0     200 100 ?
*> 30.30.30.0/24    10.10.10.2          0         100       0     200 100 ?
*> 40.40.40.0/24    0.0.0.0             0         100       32768  ?

Total number of prefixes 3
```

neighbor local-as

Use this command to specify an Autonomous System (AS) number to use with a BGP neighbor.

Use the `no` parameter with this command to disable this command.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) local-as <1-4294967295>  (no-prepend|) (replace-as|)
no neighbor (A.B.C.D|X:X::X:X|WORD) local-as <1-4294967295>
no neighbor (A.B.C.D|X:X::X:X|WORD) local-as <1-4294967295> no-prepend
no neighbor (A.B.C.D|X:X::X:X|WORD) local-as <1-4294967295> replace-as
```

For BGP unnumbered mode:

```
neighbor WORD local-as <1-4294967295>  (no-prepend|) (replace-as|)
no neighbor WORD local-as <1-4294967295>
no neighbor WORD local-as <1-4294967295> no-prepend
no neighbor WORD local-as <1-4294967295> replace-as
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in IPv4 format
`X:X::X:X`	Address of the BGP neighbor in IPv6 format
`WORD`	Name of a BGP peer group created with the neighbor WORD peer-group command. When you specify this parameter, the command applies to all peers in the group.
`<1-4294967295>`	A neighbor's AS number when extended capabilities are configured
`no-prepend`	Do not prepend local-as to update from EBGP peers
`replace-as`	Replace actual AS with local AS in the EBGP update

Note: The AS number 23456 is a reserved 2-byte AS number. An old BGP speaker (2-byte implementation) should be configured with 23456 as its remote AS number while peering with a non-mappable new BGP speaker (4-byte implementation).

Default

By default, local-as is disabled.

Command Mode

Router mode and Address Family-VRF mode and BGP unnumbered mode

Applicability

This command was introduced before OcNOS version 1.3. The new version of the command with “no-prepend” and “replace-as“ option is introduced in OcNOS version 6.4.1.

Example

The following example show a sample configuration command.

```
#configure terminal
(config)#router bgp 100
(config-router)#neighbor 20.1.1.3 remote-as 300
(config-router)#neighbor 20.1.1.3 local-as 200 no-prepend replace-as

(config)#router bgp 100
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#neighbor 3ffe:15:15:15:15::0 remote-as 300
(config-router-af)#neighbor 3ffe:15:15:15:15::0 local-as 200
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 local-as 300
```

Abbreviations

Acronym	Description
ASN	Autonomous System Number
EBGP	External Border Gateway Protocol

Port Breakout (400G) for Qumran2 Series Platforms

Overview

The port breakout capability offers a robust and secure solution for divide 400GbE ports into multiple port, ensuring a reliable network infrastructure. In today's networks, there's a demand for a diverse range of Ethernet interface speeds, including 10GbE, 25GbE, 40GbE, and 100GbE. It is essential to have a variety of cost-effective cabling options. This flexibility is crucial to address connectivity requirements and facilitate seamless migrations as network speeds and density needs continue to evolve.

Each 400GbE port (QSFP-DD) has the capacity to support up to eight SERDES, with each SERDES capable of delivering 50G of bandwidth. This capability allows for the following port configurations. The default SERDES mode operates at 50G.

Feature Characteristics

Breakout configurations facilitate the connection between network devices with varying port speeds, allowing for the optimal utilization of port bandwidth.

The breakout mode on network equipment, such as switches, routers, and servers, opens up new possibilities for network operators to keep up with the pace of bandwidth demand. By adding high-speed ports that support breakout mode, network operators can increase the front port density and incrementally enable an upgrade to higher data rates.

Benefits

The 400G platforms empower data centers and high-performance computing environments to meet the increasing demand for greater bandwidth at a reduced cost and power consumption per gigabit. Some key benefits of these platforms include:

• Upgrades from 100G to 400G systems increases the available switching bandwidth by a factor of 4, effectively addressing the need for higher data throughput.

• Enables the use of optical or copper breakouts to create higher density 100G ports, providingmore options for data connectivity and transmission.

• Reduces the number of optical fiber links, connectors, and patch panels required, achieving a fourfold reduction in infrastructure components when compared to 100G platforms with the same aggregate bandwidth. This reduction contributes to cost savings and simplifies network management.

Configuration

Use the `config# qsfp dd application` command to select the application ID to be configured for this QSFP-DD module.

Note: Only 400G application modes are supported.

Use the `show qsfp ddport no > advertisement applications` command to check the application modes.

Topology

The platform supports splitting a single 400G (QSFP-DD) port into any of the following ports.

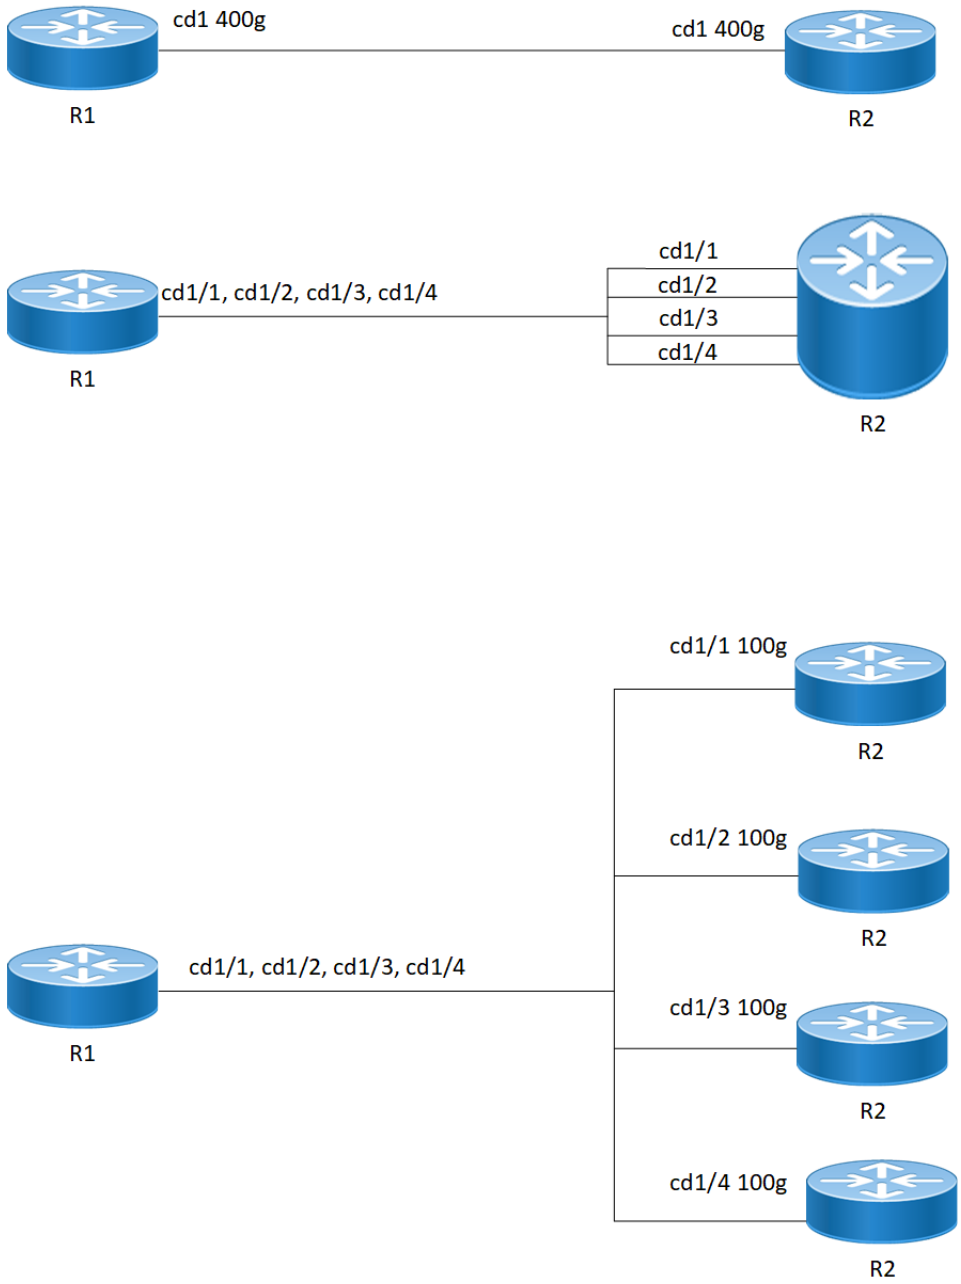

400G Port Breakout Configuration

R1

The following table outlines the configuration steps for dividing a single port into multiple ports through channelization.

`ROUTER1#configure terminal`	Enter Configuration mode.
`ROUTER1(config)# qsfp-dd 49`	Enter the QSFP-DD mode.

`ROUTER1(config-qsfp-dd)#application 3`	Select the application ID to be configured for this QSFP-DD module.
`OcNOS(config)#commit`	Commit the configuration.

EEPROM Details for ZR+ Optics

The below show command displays output for “SO-TQSFPDD4CCZRP” optics.

Execute the “`show qsfp-dd 3 eeprom`” command in the terminal window.

```
Port Number                    : 3
Identifier                     : QSFP-DD Double Density 8X Pluggable Transceiver
Name                           : SmartOptics
OUI                            : 0x0 0x53 0x4f
Part No                        : SO-TQSFPDD4CCZRP
Revision Level                 : A
Serial_Number                  : 223950575
Manufacturing Date             : 220926   (yymmddvv, v=vendor specific)
Module Power Class             : 8
Module Max Power               : 23.75 Watt
Cooling Implemented            : Yes
Module Temperature Max         : 80 Celsius
Module Temperature Min         : 0 Celsius
Operating Voltage Min          : 3.12 Volt
Optical Detector               : PIN
Rx Power Measurement           : Average Power
Tx Disable Module Wide         : No
Cable Assembly Link Length     : Separable Media
Connector Type                 : LC (Lucent Connector)
Media Interface Technology     : 1550 nm DFB
CMIS Revision                  : 4.1
Memory Model                   : Paged
MCI Max Speed                  : 1000 kHz
Active Firmware Revision       : 61.20
Inactive Firmware Revision     : 61.20
Hardware Revision              : 1.2
Media Type                     : Optical SMF
Max SMF Link Length            : 630.0 Kilometer
Wavelength Nominal             : 1547.70 nm
Wavelength Tolerance           : 166.55 nm
```

Port Breakout Configuration

Use this command to configure the port breakout on the QSFP-DD module.

R1

The following table outlines the configuration steps for port breakout.

`ROUTER1#configure terminal`	Enter Configuration mode.
`ROUTER1(config)# qsfp-dd 49`	Enter the QSFP-DD mode.
`ROUTER1(config-qsfp-dd)#application 3`	Configure the required application number. The supported range is from <2 to 15>.
`ROUTER1(config-qsfp-dd)#commit`	Commit the configuration.
`ROUTER1(config-qsfp-dd)#exit`	Exit from the QSFP-DD configuration mode.
`ROUTER1(config)#port cd49 breakout 4X100g`	Enable port breakout
`ROUTER1(config)# commit`	Commit the configuration.

Validation

Use this command to validate the port breakout configuration.

```
OcNOS#show qsfp-dd 49 application

Port Number                     : 49
-----------------------------------
   User Config    |   H/W Config
-----------------------------------
   Application 3  |   Application 3

OcNOS#show qsfp-dd 49 advertisement applications

Port Number                   : 49
> Application 1:
    | Host |
         Interface            : 400GAUI-8 C2M
         Application BR       : 425.00
         Lane Count           : 8
         Lane Sig BR          : 26.5625
         Modulation Format    : PAM4
         Bits Per Unit Intvl  : 2.000000
         Lane Assigned        : Lane-1
    | Media |
         Interface            : 400ZR, DWDM, Amplified
         Application BR       : 478.75
         Lane Count           : 1
         Lane Sig BR          : 59.84375
         Modulation Format    : DP-16QAM
         Bits Per Unit Intvl  : 8.000000
         Lane Assigned        : Lane-1
  Application 2:
```

```
  | Host |
       Interface            : 400GAUI-8 C2M
       Application BR       : 425.00
       Lane Count           : 8
       Lane Sig BR          : 26.5625
       Modulation Format    : PAM4
       Bits Per Unit Intvl  : 2.000000
       Lane Assigned        : Lane-1
  | Media |
       Interface            : 400ZR, Single Wavelen., Unamp.
       Application BR       : 478.75
       Lane Count           : 1
       Lane Sig BR          : 59.84375
       Modulation Format    : DP-16QAM
       Bits Per Unit Intvl  : 8.000000
       Lane Assigned        : Lane-1
Application 3:
  | Host |
       Interface            : 100GAUI-2 C2M
       Application BR       : 106.25
       Lane Count           : 2
       Lane Sig BR          : 26.5625
       Modulation Format    : PAM4
       Bits Per Unit Intvl  : 2.000000
       Lane Assigned        : Lane-7/Lane-5/Lane-3/Lane-1
  | Media |
       Interface            : 400ZR, DWDM, Amplified
       Application BR       : 478.75
       Lane Count           : 1
       Lane Sig BR          : 59.84375
       Modulation Format    : DP-16QAM
       Bits Per Unit Intvl  : 8.000000
       Lane Assigned        : Lane-1
Application 4:
  | Host |
       Interface            : 400GAUI-8 C2M
       Application BR       : 425.00
       Lane Count           : 8
       Lane Sig BR          : 26.5625
       Modulation Format    : PAM4
       Bits Per Unit Intvl  : 2.000000
       Lane Assigned        : Lane-1
  | Media |
       Interface            : ZR400-OFEC-16QAM
       Application BR       : 481.108374
       Lane Count           : 1
```

```
        Lane Sig BR          : 60.1385468
        Modulation Format    : DP-16QAM
        Bits Per Unit Intvl  : 8.000000
        Lane Assigned        : Lane-1
Application 5:
  | Host |
        Interface            : 100GAUI-2 C2M
        Application BR       : 106.25
        Lane Count           : 2
        Lane Sig BR          : 26.5625
        Modulation Format    : PAM4
        Bits Per Unit Intvl  : 2.000000
        Lane Assigned        : Lane-7/Lane-5/Lane-3/Lane-1
  | Media |
        Interface            : ZR400-OFEC-16QAM
        Application BR       : 481.108374
        Lane Count           : 1
        Lane Sig BR          : 60.1385468
        Modulation Format    : DP-16QAM
        Bits Per Unit Intvl  : 8.000000
        Lane Assigned        : Lane-1
Application 6:
  | Host |
        Interface            : 100GAUI-2 C2M
        Application BR       : 106.25
        Lane Count           : 2
        Lane Sig BR          : 26.5625
        Modulation Format    : PAM4
        Bits Per Unit Intvl  : 2.000000
        Lane Assigned        : Lane-7/Lane-5/Lane-3/Lane-1
  | Media |
        Interface            : ZR300-OFEC-8QAM
        Application BR       : 360.831281
        Lane Count           : 1
        Lane Sig BR          : 60.1385468
        Modulation Format    : DP-8QAM
        Bits Per Unit Intvl  : 6.000000
        Lane Assigned        : Lane-1
Application 7:
  | Host |
        Interface            : 100GAUI-2 C2M
        Application BR       : 106.25
        Lane Count           : 2
        Lane Sig BR          : 26.5625
        Modulation Format    : PAM4
        Bits Per Unit Intvl  : 2.000000
```

```
        Lane Assigned        : Lane-7/Lane-5/Lane-3/Lane-1
    | Media |
        Interface            : ZR200-OFEC-QPSK
        Application BR       : 240.554187
        Lane Count           : 1
        Lane Sig BR          : 60.1385468
        Modulation Format    : DP-QPSK
        Bits Per Unit Intvl  : 4.000000
        Lane Assigned        : Lane-1
  Application 8:
    | Host |
        Interface            : 100GAUI-2 C2M
        Application BR       : 106.25
        Lane Count           : 2
        Lane Sig BR          : 26.5625
        Modulation Format    : PAM4
        Bits Per Unit Intvl  : 2.000000
        Lane Assigned        : Lane-7/Lane-5/Lane-3/Lane-1
    | Media |
        Interface            : ZR100-OFEC-QPSK
        Application BR       : 120.277094
        Lane Count           : 1
        Lane Sig BR          : 30.069273
        Modulation Format    : DP-QPSK
        Bits Per Unit Intvl  : 4.000000
        Lane Assigned        : Lane-1
```

Port Breakout Interfaces

Use this command to configure the to see the interfaces after the port breakout.

```
ROUTER1#show interface brief | include cd49
cd49/1        ETH           --    routed           up      none    100g  --       No  No
cd49/2        ETH           --    routed           up      none    100g  --       No  No
cd49/3        ETH           --    routed           up      none    100g  --       No  No
cd49/4        ETH           --    routed           up      none    100g  --       No  No
```

Port Breakout Unconfiguration

Use this command to unconfigure the port breakout on the QSFP-DD module.

R1

The following table outlines the configuration steps for port breakout.

`ROUTER1#configure terminal`	Enter Configuration mode.
`ROUTER1(config)# qsfp-dd 49`	Enter the QSFP-DD mode.
`ROUTER1(config-qsfp-dd)#no application`	Remove the application.
`ROUTER1(config-qsfp-dd)#commit`	Commit the configuration.
`ROUTER1(config-qsfp-dd)#exit`	Exit from the QSFP-DD configuration mode.
`ROUTER1(config)#no port cd49 breakout`	Remove the port breakout. Your port will revert to functioning as a 400G port.
`ROUTER1(config)# commit`	Commit the configuration.

```
OcNOS#show qsfp-dd 49 application

Port Number                     : 49
-----------------------------------
   User Config    |   H/W Config
-----------------------------------
   Application 1  |   Application 1

ROUTER1#show interface brief | include cd49
cd49        ETH           --    routed            up      none    400g  --       No  No
```

Port Breakout Configuration with serdes 25g

Use this command to configure the port breakout on the QSFP-DD module.

R1

The following table outlines the configuration steps for port breakout.

`ROUTER1#configure terminal`	Enter Configuration mode.
`ROUTER1(config)# qsfp-dd 49`	Enter the QSFP-DD mode.
`ROUTER1(config-qsfp-dd)#application 12`	Configure the required application number. The accepted range is from 2 to 15.
`ROUTER1(config-qsfp-dd)#commit`	Commit the configuration.
`ROUTER1(config-qsfp-dd)#exit`	Exit from the QSFP-DD configuration mode.
`ROUTER1(config)#port cd49 breakout 2X100g serdes 25g`	Configure port breakout with 25G Serdes.
`ROUTER1(config)# commit`	Commit the configuration.

Validation

Use this command to validate the port breakout configuration.

```
OcNOS#show qsfp-dd 49 application
```

```
Port Number                    : 49
----------------------------------
   User Config    |   H/W Config
----------------------------------
   Application 12  |   Application 12
```

Port Breakout Interfaces

Use this command to configure the to see the interfaces after the port breakout.

```
ROUTER1#show interface brief | include cd49
cd49/1        ETH           --    routed           up      none    100g  --       No  No
cd49/2        ETH           --    routed           up      none    100g  --       No  No
```

Port Breakout Unconfiguration with serdes 25g

Use this command to unconfigure the port breakout on the QSFP-DD module.

R1

The following table outlines the configuration steps for port breakout.

`ROUTER1#configure terminal`	Enter Configuration mode.
`ROUTER1(config)# qsfp-dd 49`	Enter the QSFP-DD mode.
`ROUTER1(config-qsfp-dd)#no application`	Remove the application
`ROUTER1(config-qsfp-dd)#commit`	Commit the configuration.
`ROUTER1(config-qsfp-dd)#exit`	Exit from the QSFP-DD configuration mode.
`ROUTER1(config)#no port cd49 breakout`	Remove the port breakout. Your port will revert to functioning as a 400G port.
`ROUTER1(config)# commit`	Commit the configuration.

```
OcNOS#show qsfp-dd 49 application

Port Number                    : 49
----------------------------------
   User Config    |   H/W Config
----------------------------------
   Application 1  |   Application 1

ROUTER1#show interface brief | include cd49
cd49        ETH           --    routed           up      none    400g  --       No  No
```

Support IGMP Snooping for Provider Bridge

Overview

In Layer-2 switches, multicast IP traffic is handled in the same manner as broadcast traffic and forwards frames received on one interface to all other interfaces. This creates excessive traffic on the network, and affects network performance. The Internet Group Management Protocol (IGMP) Snooping allows switches to monitor network traffic, and determine hosts to receive multicast traffic. Thus, at a time only an host's membership report is relayed from a group instead of a report from each host in the group.

A Provider Bridge (PB) network is a virtual bridge Local Area Network (LAN) that comprises of Service provider bridges (SVLAN and PB) and attached LANs controlled under a single service provider administration. Provider bridges interconnect the MACs of the IEEE 802 LANs separately. This combined provider bridged network relay frames to all the connected LANs that provide customer interfaces for each service instance.

Feature Characteristics

The existing IGMP Snooping extended to support in the Provider Bridged (PB) network. The PB connects customer LANs using the switched provider network consisting of SVLAN bridges and provider edge bridges. Each customer LAN is connected to a separate service VLAN inside the provider network. Current release supports the IGMPv1/ IGMPv2/IGMPv3.

The following are supported:

- Snooping entries are captured in provider bridge network
- Egress traffic from router is tagged with single SVLAN ID
- IGMP snooping feature supported only in SVLAN

Benefits

This feature enables a Provider bridging network service provider to conserve bandwidth by efficiently switching the multicast packets.

Prerequisites

IGMP snooping is available over a number of network underlays. In this chapter, it is assumed that Provider Bridge support is configured.

Configuration

Topology

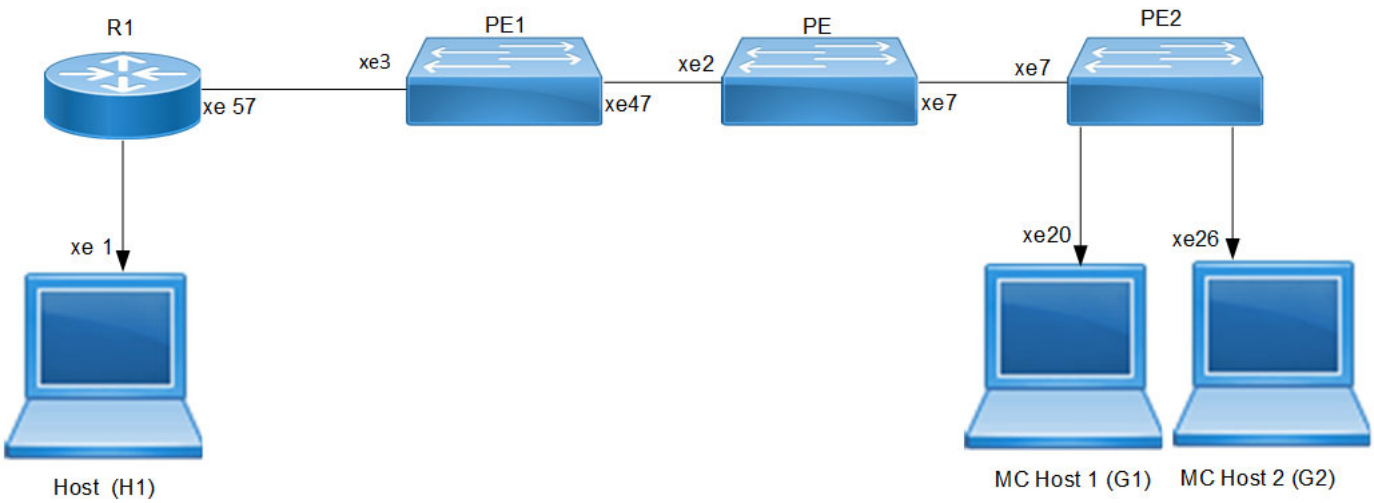

IGMP Snooping Provider Bridge Topology

R1

`#configure terminal`	Enter the configure mode.
`R1(config)#bridge 1 protocol provider-rstp edge`	Configure bridge `1` to the spanning-tree table.
`R1(config)#vlan database`	Configure the VLAN database.
`R1(config)#vlan 2 type service point-point bridge 1 state enable`	Configure the SVLAN 2 to bridge 1.
`R1(config)#ip multicast-routing`	Configure the multicast routing on the router.
`R1(config)#ip pim rp-address 1.1.1.1`	Configure Rendezvous Point (RP) address for multicast groups.
`R1(config)#interface lo`	Enter into lo interface.
`R1(config-if)#ip address 1.1.1.1/24 secondary`	Configure rp address as secondary.
`R1(config-if)#ip pim sparse-mode`	Enable the PIM sparse mode.
`R1(config-if)#exit`	Exit the loopback interface mode.
`R1(config)#interface svlan1.2`	Create the SVLAN interface.
`R1(config-if)#ip address 20.1.1.1/24`	Configure IPv4 address to VLAN interface.
`R1(config-if)#ip pim sparse-mode`	Configure PIM sparse mode.
`R1(config-if)#exit`	Exit the SVLAN interface mode.
`R1(config)#interface xe1`	Enter interface mode.
`R1(config-if)#ip address 10.1.1.1/24`	Configure IPv4 address to interface
`R1(config-if)#ip pim sparse-mode`	Configure PIM sparse mode.
`R1(config-if)#commit`	Commit the configurations.
`R1(config-if)#exit`	Exit the interface mode.
`R1(config)#interface xe57`	Enter interface mode.
`R1(config-if)#switchport`	Configure switchport.
`R1(config-if)#dot1ad ethertype 0x8100`	Configure ether type 0x8100.
`R1(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group.
`R1(config-if)#switchport mode provider-network`	Configure switchport trunk mode.
`R1(config-if)#switchport provider-network allowed vlan add 2`	Configure the VLAN to switchport trunk mode.
`R1(config-if)#commit`	Commit configurations

PE1

`#configure terminal`	Enter the configure mode.
`PE1(config)#bridge 1 protocol provider-rstp edge`	Configure bridge `1` to the spanning-tree table.
`PE1(config)#vlan database`	Configure the VLAN database.
`PE1(config)#vlan 2 type service point-point bridge 1 state enable`	Configure the SVLAN 2 to bridge 1.
`PE1(config)#ip multicast-routing`	Configure the multicast routing on the router.
`PE1(config)#interface svlan1.2`	Create VLAN interface.
`PE1(config-if)#igmp snooping enable`	Configure IPv4 address to VLAN interface .
`PE1PE1(config-if)#exit`	Exit the interface mode.
`PE1(config)#interface xe3`	Enter interface mode.
`PE1(config-if)#switchport`	Configure Switchport.
`PE1(config-if)#dot1ad ethertype 0x8100`	Configure ethertype.
`PE1(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group spanning tree disable .
`PE1(config-if)#switchport mode provider-network`	Configure provider network .
`PE1(config-if)#switchport provider-network allowed vlan add 2`	Configure the SVLAN to interface .
`PE1(config-if)#commit`	Commit configurations.
`PE1(config-if)#exit`	Exit the interface mode.
`PE1(config)#interface xe47`	Enter interface mode.
`PE1(config-if)#switchport`	Configure switchport
`PE1(config-if)#dot1ad ethertype 0x8100`	Configure ethertype.
`PE1(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group spanning tree disable.
`PE1(config-if)#switchport mode provider-network`	Configure provider network.
`PE1(config-if)#switchport provider-network allowed vlan add 2`	Configure service vlan to provider network.
`PE1(config-if)#commit`	Commit configurations.
`PE1(config-if)#exit`	Exit the interface mode.

PE

`#configure terminal`	Enter the configure mode.
`PE(config)#bridge 1 protocol provider-rstp edge`	Configure bridge `1` to the spanning-tree table.
`PE(config)#vlan database`	Configure the VLAN database
`PE(config)#vlan 2 type service point-point bridge 1 state enable`	Configure the SVLAN 2 to bridge 1.
`PE(config)#ip multicast-routing`	Configure the multicast routing on the router.
`PE(config)#interface svlan1.2`	Create VLAN interface.
`PE(config-if)#igmp snooping enable`	Configure IPv4 address to VLAN interface.
`PE(config-if)#exit`	Exit the interface mode.
`PE(config)#interface xe2`	Enter interface mode.
`PE(config-if)#switchport`	Configure Switchport
`PE(config-if)#dot1ad ethertype 0x8100`	Configure ethertype
`PE(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group spanning tree disable.
`PE(config-if)#switchport mode provider-network`	Configure provider network.
`PE(config-if)#switchport provider-network allowed vlan add 2`	Configure the SVLAN to interface.
`PE(config-if)#commit`	Commit configurations.
`PE(config-if)#exit`	Exit the interface mode.
`PE(config)#interface xe7`	Enter interface mode.
`PE(config-if)#switchport`	Configure switchport.
`PE(config-if)#dot1ad ethertype 0x8100`	Configure ethertype.
`PE(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group spanning tree disable.
`PE(config-if)#switchport mode provider-network`	Configure provider network.
`PE(config-if)#switchport provider-network allowed vlan add 2`	Configure service vlan to provider network.
`PE(config-if)#commit`	Commit configurations.
`PE(config-if)#exit`	Exit the interface mode.

PE2

`#configure terminal`	Enter the configure mode.
`PE2(config)#bridge 1 protocol provider-rstp edge`	Configure bridge `1` to the spanning-tree table.
`PE2(config)#vlan database`	Configure the VLAN database.
`PE2(config)#vlan 2 type service point-point bridge 1 state enable`	Configure the SVLAN 2 to bridge 1.
`PE2(config)#ip multicast-routing`	Configure the multicast routing on the router.
`PE2(config)#interface svlan1.2`	Create VLAN interface.
`PE2(config-if)#igmp snooping enable`	Enable the IGMP snooping on VLAN interface.
`PE2(config-if)#exit`	Exit the VLAN interface mode.
`PE2(config)#interface xe7`	Enter interface mode.
`PE2(config-if)#switchport`	Configure Switchport.
`PE2(config-if)#dot1ad ethertype 0x8100`	Configure ethertype.
`PE2(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group spanning tree disable.
`PE2(config-if)#switchport mode provider-network`	Configure provider network.
`PE2(config-if)#switchport provider-network allowed vlan add 2`	Configure the SVLAN to interface.
`PE2(config-if)#commit`	Commit configurations.
`PE2(config-if)#exit`	Exit the interface mode.
`PE2(config)#interface xe20`	Enter interface mode.
`PE2(config-if)#switchport`	Configure switchport.
`PE2(config-if)#dot1ad ethertype 0x8100`	Configure ethertype.
`PE2(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group spanning tree disable.
`PE2(config-if)#switchport mode provider-network`	Configure provider network.
`PE2(config-if)#switchport provider-network allowed vlan add 2`	Configure service VLAN to provider network.
`PE2(config-if)#commit`	Commit configurations.
`PE2(config-if)#exit`	Exit the interface mode.
`PE2(config)#interface xe22`	Enter interface mode.
`PE2(config-if)#switchport`	Configure switchport.
`PE2(config-if)#dot1ad ethertype 0x8100`	Configure ethertype.
`PE2(config-if)#bridge-group 1 spanning-tree disable`	Configure bridge group spanning tree disable.

`PE2(config-if)#switchport mode provider-network`	Configure provider network.
`PE2(config-if)#switchport provider-network allowed vlan add 2`	Configure service VLAN to provider network.
`PE2(config-if)#commit`	Commit configurations.
`PE2(config-if)#exit`	Exit the interface mode.

Validation

R1

```
MCRTR#show ip igmp groups
IGMP Instance wide G-Recs Count is: 2
IGMP Connected Group Membership
Group Address    Interface            Uptime     Expires   State    Last Reporter
231.1.1.1        svlan1.2             00:00:12   00:04:07  Active   0.0.0.0
231.1.1.2        svlan1.2             00:00:12   00:04:07  Active   0.0.0.0
MCRTR#
MCRTR#show ip pim mroute
IP Multicast Routing Table

(*,*,RP) Entries: 0
G/prefix Entries: 0
(*,G) Entries: 2
(S,G) Entries: 0
(S,G,rpt) Entries: 0
FCR Entries: 0

(*, 231.1.1.1)
RP: 1.1.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: JOINED
 Local     ..i.............................
 Joined    ................................
 Asserted  ................................
FCR:

(*, 231.1.1.2)
RP: 1.1.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: JOINED
 Local     ..i.............................
 Joined    ................................
 Asserted  ................................
FCR:

MCRTR#
```

PE1

```
PEB1-7014#show igmp snooping interface
Global IGMP Snooping information
 IGMP Snooping Enabled
 IGMPv1/v2 Report suppression Enabled
 IGMPv3 Report suppression Enabled

IGMP Snooping information for svlan1.2
 IGMP Snooping enabled
 Snooping Querier none
 IGMP Snooping other querier timeout is 255 seconds
 Group Membership interval is 260 seconds
 IGMPv2 fast-leave is disabled
 IGMPv1/v2 Report suppression enabled
 IGMPv3 Report suppression enabled
 Router port detection using IGMP Queries
 Number of router-ports: 1
 Number of Groups: 0
 Number of v1-reports: 0
 Number of v2-reports: 0
 Number of v2-leaves: 0
 Number of v3-reports: 0
 Active Ports:
  xe3
  xe47
PEB1-7014#show igmp snooping groups
IGMP Instance wide G-Recs Count is: 2
IGMP Snooping Group Membership
Group source list: (R - Remote, S - Static, > - Hw Installed)
Vlan   Group/Source Address    Interface      Flags   Uptime      Expires  Last Reporter     Version
2      231.1.1.1              xe47           R    >  00:07:15    00:03:48 0.0.0.0           V3
2      231.1.1.2              xe47           R    >  00:07:15    00:03:48 0.0.0.0           V3
PEB1-7014#
```

PE

```
PB-7024#show igmp snooping interface
Global IGMP Snooping information
 IGMP Snooping Enabled
 IGMPv1/v2 Report suppression Enabled
 IGMPv3 Report suppression Enabled

IGMP Snooping information for svlan1.2
 IGMP Snooping enabled
 Snooping Querier none
 IGMP Snooping other querier timeout is 255 seconds
 Group Membership interval is 260 seconds
 IGMPv2 fast-leave is disabled
 IGMPv1/v2 Report suppression enabled
 IGMPv3 Report suppression enabled
```

```
 Router port detection using IGMP Queries
 Number of router-ports: 1
 Number of Groups: 0
 Number of v1-reports: 0
 Number of v2-reports: 0
 Number of v2-leaves: 0
 Number of v3-reports: 0
 Active Ports:
  xe7
  xe2
PB-7024#
PB-7024#show igmp snooping groups
IGMP Instance wide G-Recs Count is: 2
IGMP Snooping Group Membership
Group source list: (R - Remote, S - Static, > - Hw Installed)
Vlan   Group/Source Address    Interface     Flags   Uptime
Expires  Last Reporter     Version
2      231.1.1.1              xe7           R    >  00:07:15        00:03:45 20.1.1.2          V3
2      231.1.1.2              xe7           R    >  00:07:15        00:03:51 20.1.1.3          V3
PB-7024#
```

PE2

```
PEB2-7019#show igmp snooping interface
Global IGMP Snooping information
 IGMP Snooping Enabled
 IGMPv1/v2 Report suppression Disabled
 IGMPv3 Report suppression Disabled

IGMP Snooping information for svlan1.2
 IGMP Snooping enabled
 Snooping Querier none
 IGMP Snooping other querier timeout is 255 seconds
 Group Membership interval is 260 seconds
 IGMPv2 fast-leave is disabled
 IGMPv1/v2 Report suppression disabled
 IGMPv3 Report suppression disabled
 Router port detection using IGMP Queries
 Number of router-ports: 1
 Number of Groups: 0
 Number of v1-reports: 0
 Number of v2-reports: 0
 Number of v2-leaves: 0
 Number of v3-reports: 0
 Active Ports:
  xe20
  xe26
  xe7
    PEB2-7019#
    PEB2-7019#show igmp snooping groups
    IGMP Instance wide G-Recs Count is: 2
    IGMP Snooping Group Membership
```

```
    Group source list: (R - Remote, S - Static, > - Hw Installed)
Vlan  Group/Source Address   Interface    Flags  Uptime  Expires Last Reporter   Version
2      231.1.1.1              xe20          R    >  00:07:14  00:03:45 20.1.1.2          V3
2      231.1.1.2              xe26          R    >  00:07:15  00:03:51 20.1.1.3       V3
PEB2-7019#
```

Abbreviations

Acronym	Description
IGMP	Internet Group Management Protocol
PB	Provider Bridged
SVLAN	Service Provider VLAN

TCP MSS configuration for BGP neighbors

Overview

The manual configuration between the routing devices establishes the BGP peer that creates a TCP session.

This feature enables the configuration of TCP Maximum Segment Size (MSS) that defines the maximum segment size in a single TCP segment during a communication session. TCP segment is a unit of data transmitted in a TCP connection. TCP uses three-way handshake process for initial establishment of a TCP connection. In the three-way handshake process, the sending host sends a SYN packet. Once the receiving host receives the SYN packet, it acknowledges and sends back a SYN-ACK packet to the sending host. Once the sending host receives the SYN-ACK packet from the receiving host, it sends an ACK packet, establishing a reliable connection. In this three way handshake process, the MSS is negotiated between the BGP neighbors.

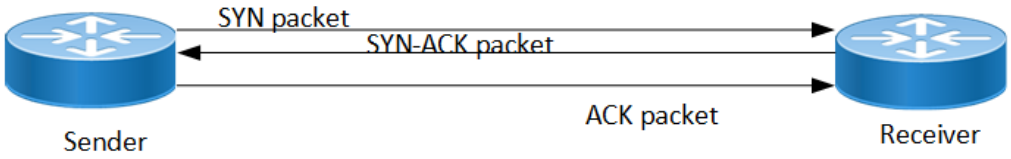

Three-way handshake

Feature Characteristics

The configuration of the TCP MSS for BGP neighbors helps the neighbors adjust the MSS value of the TCP SYN packet. Configure the TCP MSS through the CLI and NetConf interface. The configurable MSS range is offered from 40-1440 bytes. By default, the MTU value for ethernet cable is 1500 bytes. When configuring the highest MSS value that is 1440, the total MSS becomes 1440 bytes (MSS) plus 20 bytes (IP Header Size), 20 bytes (TCP Header), and Ethernet header which does not cross the default path MTU value.

TCP MSS for BGP neighbor

Benefits

By default, the interface MTU value determines the MSS value of a packet. When the interface MTU value exceeds the default ethernet path MTU value of 1500 bytes, the MSS value also crosses the default ethernet path MTU value, resulting in packet fragmentation. The configuration of the specific MSS value limits the packet size irrespective of the interface MTU value, preventing packet fragmentation.

Prerequisites

Requires the knowledge on TCP handshake and BGP neighbor discovery.

Configuration

This section shows the procedure to configure TCP MSS between BGP peers.

Topology

The below example shows the configuration required to enable BGP on an interface. PE1 and RR1 are routers belonging to the same Autonomous System (AS) with the Autonomous System Number (ASN) as AS100, connecting to network 10.1.1.0/24. First, define the routing process and the ASN to which the routers belong. Then, define BGP neighbors to start exchanging routing updates and configure the TCP MSS for BGP between PE1 and RR1 devices.

TCP MSS for BGP neighbor

Configuration

The configuration shows how to configure the TCP MSS value for the BGP peer.

PE1

`PE1#configure terminal`	Enter Configuration mode.
`PE1(config)#interface lo`	Enter interface mode for loopback.
`PE1(config-if)#ip address 1.1.1.1/32 secondary`	Specify the interface IP address `1.1.1.1`.
`PE1(config-if)#exit`	Exit the interface mode.
`PE1(config)#interface xe1`	Enter interface mode for `xe1`.
`PE1(config-if)#ip address 10.1.1.1/24`	Specify the IP address `10.1.1.1` for the interface.
`PE1(config-if)#exit`	Exit interface mode for `xe1`.
`PE1(config)#router bgp 100`	Define the routing process. The number 100 specifies the ASN of PE1.
`PE1(config-router)#bgp router-id 1.1.1.1`	Configure bgp router-id same as loopback IP address`1.1.1.1`.
`PE1(config-router)#neighbor 10.1.1.2 remoteas 100`	Define BGP neighbors, and establish a TCP session. `10.1.1.2` is the IP address of the neighbor and 100 is the neighbor's ASN.

`PE1(config-router)#neighbor 10.1.1.2 tcp-mss 800`	Configure TCP MSS value.
`PE1(config-router)#address-family ipv4 unicast`	Enter address-family IPv4 unicast mode.
`PE1(config-router-af)#neighbor 10.1.1.2 activate`	Activate neighbor with IP address `10.1.1.2` in the IPv4 address family.
`PE1(config-router-af)#redistribute connected`	Redistributing connected routes inside BGP.
`PE1(config-router-af)#exit-address-family`	Exit address-family mode.
`PE1(config-router)#commit`	Commit the candidate configuration to the running configuration.

RR1

`RR1#configure terminal`	Enter configuration mode.
`RR1(config)#interface lo`	Enter interface mode for loopback.
`RR1(config-if)#ip address 2.2.2.2/32 secondary`	Specify the interface address `2.2.2.2`.
`RR1(config-if)#exit`	Exit interface mode.
`RR1(config)#interface xe47`	Enter interface mode for `xe47`.
`RR1(config-if)#ip address 10.1.1.2/24`	Specify IP address `10.1.1.2/24` for the interface.
`RR1(config-if)#exit`	Exit interface mode for `xe47`.
`RR1(config)#router bgp 100`	Define the routing process. The number 100 specifies the ASN of RR1.
`RR1(config-router)#bgp router-id 2.2.2.2`	Configure BGP router-id same as loopback IP address `2.2.2.2`.
`RR1(config-router)#neighbor 10.1.1.1 remoteas 100`	Define BGP neighbors, and establish a TCP session. `10.1.1.1` is the ip address of the neighbor and 100 is the neighbor's ASN.
`RR1(config-router)#neighbor 10.1.1.1 passive`	Configure BGP neighbor `10.1.1.1` passive.
`RR1(config-router)#address-family ipv4 unicast`	Enter address-family IPv4 unicast mode
`RR1(config-router-af)#neighbor 10.1.1.1 activate`	Activate the neighbor in the IPv4 address family.
`RR1(config-router-af)#neighbor 10.1.1.1 route-reflector-client`	Configure RR1 as the Route-Reflector (RR) and neighbor PE1 as its client.
`RR1(config-router-af)#redistribute connected`	Redistributing connected routes inside BGP.
`RR1(config-router-af)#exit-address-family`	Exit address-family mode.
`RR1(config-router)#commit`	Commit the candidate configuration to the running configuration.

Validation

PE1

```
PE1#show bgp summary
BGP router identifier 1.1.1.1, local AS number 100
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Dow
n   State/PfxRcd
10.1.1.2                 4   100  171        170       1      0       0  00:00:11
               0

Total number of neighbors 1

Total number of Established sessions 1
PE1#

PE1#show bgp neighbors
BGP neighbor is 10.1.1.2, remote AS 100, local AS 100, internal link, peer index
: 2
  BGP version 4, local router ID 10.1.1.1, remote router ID 10.1.1.2
  BGP state = Established, up for 00:07:29
  Last read 00:00:24, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 43 messages, 1 notifications, 0 in queue
  Sent 46 messages, 4 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds

For address family: IPv4 Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  0 announced prefixes

 Connections established 6; dropped 5
Local host: 10.1.1.1, Local port: 34738
Foreign host: 10.1.1.2, Foreign port: 179
TCP MSS: (800), Advertise TCP MSS: (800), Send TCP MSS: (800),  Receive TCP MSS:
 (536)
Sock FD : (25)
Nexthop: 10.1.1.1
Nexthop global: ::
```

```
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:08:45, due to Administratively Reset (Cease Notification sent)
```

RR1

```
RR1#show bgp summary
BGP router identifier 2.2..2.2, local AS number 100
BGP table version is 1
0 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/Dow
n   State/PfxRcd
10.1.1.1                 4   100    2          3       1      0       0  00:00:26
               0

Total number of neighbors 1

Total number of Established sessions 1

RR1#show bgp neighbors
BGP neighbor is 10.1.1.1, remote AS 100, local AS 100, internal link, peer index
: 2
  BGP version 4, local router ID 10.1.1.2, remote router ID 10.1.1.1
  BGP state = Established, up for 00:08:31
  Last read 00:00:24, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
  Received 46 messages, 4 notifications, 0 in queue
  Sent 47 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds

For address family: IPv4 Unicast  BGP table version 1, neighbor version 1
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  0 accepted prefixes
  0 announced prefixes

 Connections established 6; dropped 5
Local host: 10.1.1.2, Local port: 179
Foreign host: 10.1.1.1, Foreign port: 34738
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (800),  Receive TCP MSS:
(536)
Sock FD : (22)
Nexthop: 10.1.1.2
Nexthop global: ::
```

```
Nexthop local: ::
BGP connection: non shared network
Last Reset: 00:09:52, due to BGP Notification received
```

New CLI Commands

neighbor tcp-mss

Use this command to set the BGP TCP MSS of a neighbor.

Use the `no` parameter with this command to remove a TCP MSS setting from a BGP neighbor.

Command Syntax

```
neighbor (A.B.C.D|X:X::X:X|WORD) tcp-mss <40-1440>
no neighbor (A.B.C.D|X:X::X:X|WORD) tcp-mss
```

For BGP unnumbered mode:

```
neighbor WORD tcp-mss <40-1440>
no neighbor WORD tcp-mss
```

Parameters

`A.B.C.D`	Address of the BGP neighbor in an IPv4 format
`X:X::X:X`	Address of the BGP neighbor in an IPv6 format
`WORD`	Name of a BGP peer group created with the *neighbor WORD peer-group* command. When you specify this parameter, the command applies to all peers in the group.
`<40-1440>`	Configure TCP MSS

Default

By default, `neighbor tcp-mss` is disabled.

Command Mode

Router mode, address family-vrf mode and BGP unnumbered mode.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
#configure terminal
(config)#router bgp 10
(config-router)#neighbor 10.10.0.72 tcp-mss 1000
(config)#router bgp 100
(config-router)#address-family ipv6 vrf VRF_A
(config-router-af)#neighbor 3ffe:15:15:15:15::0 tcp-mss 900
```

For unnumbered peer below configuration is given in BGP unnumbered-mode.

```
(config)#router bgp 100
```

```
(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor eth1 tcp-mss 800
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
ACK	Acknowledgment
BGP	Border Gateway Protocol
TCP	Transmission Control Protocol
MSS	Maximum Segment Size
MTU	Maximum Transmission Unit
SYN	Synchronize

Glossary

The following provides definitions for key terms used throughout this document.

BGP	BGP is an exterior gateway protocol to exchange route information and interconnect various networks on the global internet.
BGP neighbor	BGP neighbors, called peers, are established by manual configuration among routers to create a TCP session on port 179, which exchanges routing information between two systems, defined by their Autonomous System Numbers (ASNs).
MSS	MSS is a TCP parameter that defines the maximum amount of data in a TCP segment that can be transmitted. TCP - TCP is one of the main protocols in the Internet Protocol (IP) suite. It offers a secure and reliable connection between two devices.
TCP	TCP is one of the main protocols in the Internet Protocol (IP) suite. It offers a secure and reliable connection between two devices.
TCP segment	TCP segment is a unit of data transmitted in a TCP connection. The segment consists of header and payload. The header contains the control information to manage the transmission, and the payload contains the actual data that needs to be transmitted.

TCP MSS configuration for LDP sessions

Overview

Label Distribution Protocol (LDP) uses Transmission Control Protocol (TCP) to establish sessions between the devices. This feature enables the configuration of TCP Maximum Segment Size (MSS) that defines the maximum segment size in a single TCP segment during a communication session.TCP segment is a unit of data transmitted in a TCP connection. TCP uses three-way handshake process for initial establishment of a TCP connection. In the three-way handshake process, the sending host sends a SYN packet. Once the receiving host receives the SYN packet, it acknowledges and sends back a SYN-ACK packet to the sending host. Once the sending host receives the SYN-ACK packet from the receiving host, it sends an ACK packet, establishing a reliable connection. In this three way handshake process, the MSS is negotiated between the LDP neighbors.

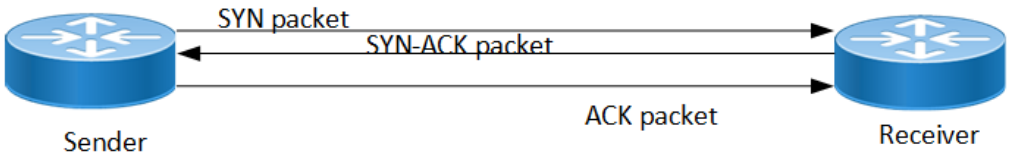

Three-way handshake

Feature Characteristics

The configuration of the TCP MSS for LDP neighbors helps the neighbors adjust the MSS value of the TCP SYN packet. Configure the TCP MSS through the CLI and NetConf interface. The configurable MSS range is offered from 560 to 1440. By default, the MTU value for ethernet cable is 1500 bytes. When configuring the highest MSS value that is 1440, the total MSS becomes 1440 bytes (MSS) plus 20 bytes (IP Header Size), 20 bytes (TCP Header), and Ethernet header which does not cross the default path MTU value.

Note: After configuring TCP MSS, use *clear ldp session* command to apply the MSS for the operational session.

Configuring TCP MSS

Benefits

By default, the interface MTU value determines the MSS value of an LDP packet. When the interface MTU value exceeds the default ethernet path MTU value of 1500 bytes, the MSS value also crosses the default ethernet path MTU

value, resulting in packet fragmentation. The configuration of the specific MSS value limits the packet size irrespective of the interface MTU value, preventing packet fragmentation.[

Prerequisites

Requires the knowledge on TCP handshake and the formation of LDP neighbors.

Configuration

This section shows the procedure to configure TCP MSS for LDP session.

Enable Label Switching

Running LDP on a system requires the following tasks:

1. Enabling label-switching on the interface on NSM.
2. Enabling LDP on an interface in the LDP daemon.
3. Running an Internal Gateway Protocol (IGP), for example, Open Shortest Path first (OSPF), to distribute reachability information within the MPLS cloud.
4. Configuring the transport address.
5. Configure the TCP MSS neighbor on peer node (Active node).

Topology

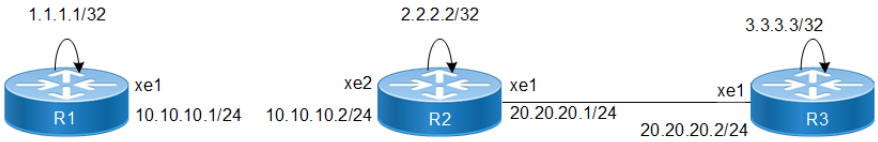

Device topology for TCP MSS for LDP

Configuration

The below configuration shows how to configure the TCP MSS value for the LDP neighbors.

R1 - NSM

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R1(config-if)#ip address 10.10.10.1/24`	Assign IP address `10.10.10.1/24` to interface.
`R1(config-if)#label-switching`	Enable label switching on interface `xe1`.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface lo`	Specify the loopback interface to be configured.

`R1(config-if)#ip address 1.1.1.1/32 secondary`	Set the IP address of the loopback interface to `1.1.1.1/32`.
`R1(config-if)#commit`	Commit the transaction.

R1 - LDP

`R1(config)#router ldp`	Enter Router mode for LDP.
`R1(config-router)#router-id 1.1.1.1`	Set the router ID to IP address `1.1.1.1`.
`R1(config-router)#transport-address ipv4 1.1.1.1`	Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter `ipv6` if you are configuring an IPv6 interface.
`R1(config-router)#targeted-peer ipv4 3.3.3.3`	Configure targeted peer `3.3.3.3`.
`R1(config-router-targeted-peer)#exit`	Exit targeted peer-mode.
`R1(config-router)#exit`	Exit the router mode and return to the configure mode.
`R1(config)#interface xe1`	Enter interface mode `xe1`.
`R1(config-if)#enable-ldp ipv4`	Enable LDP on `xe1`.
`R1(config-if)#commit`	Commit the transaction.

R1 - OSPF

`R1(config)#router ospf 100`	Configure the routing process and specify the process ID `100`. The process ID should be a unique positive integer identifying the routing process.
`R1(config-router)#network 10.10.10.0/24 area 0`	Define the interface `10.10.10.0/24`,on which OSPF runs and associate the area ID `0` with the interface.
`R1(config-router)#network 1.1.1.1/32 area 0`	Define the interface `1.1.1.1/32`, on which OSPF runs and associate the area ID `0` with the interface.
`R1(config-router)#commit`	Commit the transaction.

R2 - NSM

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface lo`	Specify the loopback interface to be configured.
`R2(config-if)#ip address 2.2.2.2/32 secondary`	Set the IP address of the loopback interface to `2.2.2.2/32`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R2(config-if)#ip address 20.20.20.1/24`	Assign IP address `20.20.20.1/24` to interface.
`R2(config-if)#label-switching`	Enable label switching on interface `xe1`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe2`	Specify the interface `xe2` to be configured.
`R2(config-if)#ip address 10.10.10.2/24`	Assign IP address `10.10.10.2/24` to interface.

`R2(config-if)#label-switching`	Enable label switching on interface `xe2`.
`R2(config-if)#commit`	Commit the transaction.

R2 - LDP

`R2(config)#router ldp`	Enter Router mode.
`R2(config-router)#router-id 2.2.2.2`	Set the router ID to IP address `2.2.2.2`.
`R2(config-router)#transport-address ipv4 2.2.2.2`	Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter `ipv6` if you are configuring an IPv6 interface.
`R2(config-router)#neighbor 1.1.1.1 tcp-mss 600`	Configure the TCP MSS value on peer node which have active side only.
`R2(config-router)#exit`	Exit router mode and return to configure mode.
`R2(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R2(config-if)#enable-ldp ipv4`	Enable LDP on a specified interface `xe1`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe2`	Specify the interface `xe2` to be configured.
`R2(config-if)#enable-ldp ipv4`	Enable LDP on a specified interface `xe2`.
`R2(config-if)#commit`	Commit the transaction.

R2 - OSPF

`R2(config)#router ospf 100`	Configure the routing process and specify the process ID `100`. The process ID should be a unique positive integer identifying the routing process.
`R2(config-router)#network 10.10.10.0/24 area 0`	Define the interfaces `10.10.10.0/24`, on which OSPF runs and associate the area ID `0` with them.
`R2(config-router)#network 20.20.20.0/24 area 0`	Define the interfaces `20.20.20.0/24`, on which OSPF runs and associate the area ID `0` with them.
`R2(config-router)#network 2.2.2.2/32 area 0`	Define the interfaces `2.2.2.2/32`, on which OSPF runs and associate the area ID `0` with them.
`R2(config-router)#commit`	Commit the transaction.

R3 - NSM

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface lo`	Specify the loopback interface to be configured.
`R3(config-if)#ip address 3.3.3.3/32 secondary`	Set the IP address of the loopback interface to `3.3.3.3/32`.
`R3(config-if)#exit`	Exit interface mode.
`R3(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R3(config-if)#ip address 20.20.20.2/24`	Set the IP address of the interface to `20.20.20.2/24`.

`R3(config-if)#label-switching`	Enable label switching on interface `xe1`.
`R3(config-if)#commit`	Commit the transaction.

R3 - LDP

`R3(config)#router ldp`	Enter Router mode.
`R3(config-router)#router-id 3.3.3.3`	Set the router ID for IP address `3.3.3.3`.
`R3(config-router)#transport-address ipv4 3.3.3.3`	Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter `ipv6` if you are configuring an IPv6 interface.
`R3(config-router)#neighbor 2.2.2.2 tcp-mss 650`	Configure the TCP MSS value on peer node which have active side only.
`R3(config-router)#targeted-peer ipv4 1.1.1.1`	Configure targeted peer.
`R3(config-router-targeted-peer)#exit`	Exit targeted peer-mode.
`R3(config-router)#exit`	Exit the router mode and return to the configure mode.
`R3(config)#interface xe1`	Enter interface mode `xe1`.
`R3(config-if)#enable-ldp ipv4`	Enable LDP on `xe1`.
`R3(config-if)#commit`	Commit the transaction.

R3 - OSPF

`R3(config)#router ospf 100`	Configure the routing process and specify the Process ID `100`. The Process ID should be a unique positive integer identifying the routing process.
`R3(config-router)#network 20.20.20.0/24 area 0`	Define the interfaces `20.20.20.0/24`, on which OSPF runs and associate the area ID `0` with them.
`R3(config-router)#network 3.3.3.3/32 area 0`	Define the interfaces `3.3.3.3/32`, on which OSPF runs and associate the area ID `0` with them.
`R3(config-router)#commit`	Commit the transaction.

Validation

R3

```
R3#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State       KeepAlive  UpTime
       2.2.2.2            xe1        Active    OPERATIONAL   30    00:03:06
       1.1.1.1            xe1        Active    OPERATIONAL   30    00:03:06
```

```
R3#show ldp targeted-peer count
-----------------------------------------------
Num Targeted Peers: 1          [UP: 1]
-----------------------------------------------
PE2#show ldp session count
-------------------------------------------
 Multicast Peers      : 1          [UP: 1]
 Targeted Peers       : 1          [UP: 1]
 Total Sessions       : 2          [UP: 2]
-------------------------------------------
R3#show ldp routes
Prefix Addr         Nexthop Addr     Intf        Owner
1.1.1.1/32          20.20.20.1       xe1         ospf
2.2.2.2/32          20.20.20.1       xe1         ospf
3.3.3.3/32          0.0.0.0          lo          connected
10.10.10.0/24       20.20.20.1       xe1         ospf
20.20.20.0/24       0.0.0.0          xe1         connected
R3#show ldp fec-ipv4 count
------------------------
Num. IPv4 FEC(s): 5
------------------------
R3#show ldp session 2.2.2.2
Session state                 : OPERATIONAL
Session role                  : Active
TCP Connection                : Established
IP Address for TCP            : 2.2.2.2
Interface being used          : xe1
Peer LDP ID                   : 2.2.2.2:0
Preferred Peer LDP Password   : Not Set
Adjacencies                   : 20.20.20.1
Advertisement mode            : Downstream Unsolicited
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
Configured TCP MSS            : 650
Applied TCP MSS               : 650
Preferred TCP MSS             : NA
Address List received         : 2.2.2.2
                                10.10.10.2
                                20.20.20.1
Received Labels :       Fec             Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
                IPV4:10.10.10.0/24      impl-null       none
                IPV4:2.2.2.2/32         impl-null       none
                IPV4:1.1.1.1/32         25600           none
Sent Labels :   Fec                     Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
                IPV4:3.3.3.3/32         impl-null       none
```

R2

```
R2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
       3.3.3.3            xe1        Passive   OPERATIONAL   30    00:06:10
       1.1.1.1            xe2       Active    OPERATIONAL   30    00:06:10
R2#show ldp session count
-------------------------------------------
 Multicast Peers      : 2          [UP: 2]
 Targeted Peers       : 0          [UP: 0]
 Total Sessions       : 2          [UP: 2]
-------------------------------------------
R2#show ldp routes
Prefix Addr         Nexthop Addr     Intf        Owner
1.1.1.1/32          10.10.10.1       xe2        ospf
2.2.2.2/32          0.0.0.0          lo          connected
3.3.3.3/32          20.20.20.2       xe1         ospf
10.10.10.0/24       0.0.0.0          xe2        connected
20.20.20.0/24       0.0.0.0          xe1         connected

R2#show ldp session 1.1.1.1
Session state                 : OPERATIONAL
Session role                  : Active
TCP Connection                : Established
IP Address for TCP            : 1.1.1.1
Interface being used          : xe2
Peer LDP ID                   : 1.1.1.1:0
Preferred Peer LDP Password   : Not Set
Adjacencies                   : 10.10.10.1
Advertisement mode            : Downstream Unsolicited
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
Configured TCP MSS            : 600
Applied TCP MSS               : 600
Preferred TCP MSS             : NA
Address List received         : 1.1.1.1
                                10.10.10.1
                                48.48.48.48
Received Labels :       Fec             Label           Maps To
                IPV4:10.10.10.0/24      impl-null       none
                IPV4:1.1.1.1/32         impl-null       25600
Sent Labels :   Fec                     Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
                IPV4:10.10.10.0/24      impl-null       none
```

```
                IPV4:3.3.3.3/32         25601           impl-null
                IPV4:2.2.2.2/32         impl-null       none
```

R1

```
R1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
       2.2.2.2            xe1        Passive   OPERATIONAL   30    00:07:12
       3.3.3.3            xe1        Passive   OPERATIONAL   30    00:07:12
R1#show ldp session count
-----------------------------------------
 Multicast Peers      : 1          [UP: 1]
 Targeted Peers       : 1          [UP: 1]
 Total Sessions       : 2          [UP: 2]
-----------------------------------------
R1#show ldp targeted-peer count
-------------------------------------------------
Num Targeted Peers: 1          [UP: 1]
-------------------------------------------------
R1#show ldp routes
Prefix Addr         Nexthop Addr     Intf        Owner
1.1.1.1/32          0.0.0.0          lo          connected
2.2.2.2/32          10.10.10.2       xe1        ospf
3.3.3.3/32          10.10.10.2       xe1         ospf
10.10.10.0/24       0.0.0.0          xe1         connected
20.20.20.0/24       10.10.10.2       xe1         ospf
R1#show ldp fec
LSR codes    : E/N - LSR is egress/non-egress for this FEC,
               L - LSR received a label for this FEC,
               > - LSR will use this route for the FEC
FEC                 Code    Session             Out Label   ELC   Nexthop Addr
1.1.1.1/32          E >     non-existent        none        No    connected
2.2.2.2/32          NL>     2.2.2.2             impl-null   No    10.10.10.2
3.3.3.3/32          NL>     2.2.2.2             25601       No    10.10.10.2
10.10.10.0/24       NL      2.2.2.2             impl-null   No    connected
                    E >     non-existent        none        No    connected
20.20.20.0/24       NL>     2.2.2.2             impl-null   No    10.10.10.2
48.48.48.48/32      E >     non-existent        none        No    connected
```

Configure TCP MSS on ALL neighbor

R1 - NSM

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R1(config-if)#ip address 10.10.10.1/24`	Assign IP address `10.10.10.1/24` to interface.
`R1(config-if)#label-switching`	Enable label switching on interface `xe1`.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface lo`	Specify the loopback interface to be configured.
`R1(config-if)#ip address 1.1.1.1/32 secondary`	Set the IP address of the loopback interface to 1.1.1.1/32.
`R1(config-if)#commit`	Commit the transaction.

R1 - LDP

`R1(config)#router ldp`	Enter Router mode for LDP.
`R1(config-router)#router-id 1.1.1.1`	Set the router ID to IP address `1.1.1.1`.
`R1(config-router)#transport-address ipv4 1.1.1.1`	Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter `ipv6` if you are configuring an IPv6 interface.
`R1(config-router)#targeted-peer ipv4 3.3.3.3`	Configure targeted peer.
`R1(config-router)#neighbor all tcp-mss 700`	Configure the TCP MSS value with `all` neighbor.
`R1(config-router-targeted-peer)#exit`	Exit-targeted-peer-mode.
`R1(config-router)#exit`	Exit the Router mode and return to the Configure mode.
`R1(config)#interface xe1`	Enter interface mode `xe1`.
`R1(config-if)#enable-ldp ipv4`	Enable LDP on `xe1`.
`R1(config-if)#commit`	Commit the transaction.

R1 - OSPF

`R1(config)#router ospf 100`	Configure the routing process and specify the process ID (100). The process ID should be a unique positive integer identifying the routing process.
`R1(config-router)#network 10.10.10.0/24 area 0`	Define the interface `10.10.10.0/24`, on which OSPF runs and associate the area ID (0) with the interface.
`R1(config-router)#network 1.1.1.1/32 area 0`	Define the interface `1.1.1.1/32`, on which OSPF runs and associate the area ID (0) with the interface.
`R1(config-router)#commit`	Commit the transaction.

R2 - NSM

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface lo`	Specify the loopback (lo) interface to be configured.
`R2(config-if)#ip address 2.2.2.2/32 secondary`	Set the IP address of the loopback interface to `2.2.2.2/32`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R2(config-if)#ip address 20.20.20.1/24`	Assign IP address 20.20.20.1/24 to interface.
`R2(config-if)#label-switching`	Enable label switching on interface xe1.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe2`	Specify the interface `xe2` to be configured.
`R2(config-if)#ip address 10.10.10.2/24`	Assign IP address `10.10.10.2/24` to interface.
`R2(config-if)#label-switching`	Enable label switching on interface `xe2`.
`R2(config-if)#commit`	Commit the transaction.

R2 - LDP

`R2(config)#router ldp`	Enter Router mode.
`R2(config-router)#router-id 2.2.2.2`	Set the router ID to IP address `2.2.2.2`.
`R2(config-router)#transport-address ipv4 2.2.2.2`	Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter `ipv6` if you are configuring an IPv6 interface.
`R2(config-router)#neighbor all tcp-mss 710`	Configure the TCP MSS value with `all` neighbor.
`R2(config-router)#exit`	Exit Router mode and return to configure mode.
`R2(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R2(config-if)#enable-ldp ipv4`	Enable LDP on a specified interface `xe1`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe2`	Specify the interface `xe2` to be configured.
`R2(config-if)#enable-ldp ipv4`	Enable LDP on a specified interface `xe2`.
`R2(config-if)#commit`	Commit the transaction.

R2 - OSPF

`R2(config)#router ospf 100`	Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process.
`R2(config-router)#network 10.10.10.0/24 area 0`	Define the interfaces `10.10.10.0/24`, on which OSPF runs and associate the area ID (0) with them.
`R2(config-router)#network 20.20.20.0/24 area 0`	Define the interfaces `20.20.20.0/24`, on which OSPF runs and associate the area ID (0) with them.

`R2(config-router)#network 2.2.2.2/32 area 0`	Define the interfaces `2.2.2.2/32`, on which OSPF runs and associate the area ID (0) with them.
`R2(config-router)#commit`	Commit the transaction.

R3 - NSM

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface lo`	Specify the loopback interface to be configured.
`R3(config-if)#ip address 3.3.3.3/32 secondary`	Set the IP address of the loopback interface to `3.3.3.3/32`.
`R3(config-if)#exit`	Exit interface mode.
`R3(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R3(config-if)#ip address 20.20.20.2/24`	Set the IP address of the interface to `20.20.20.2/24`.
`R3(config-if)#label-switching`	Enable label switching on interface `xe1`.
`R3(config-if)#commit`	Commit the transaction.

R3 - LDP

`R3(config)#router ldp`	Enter Router mode.
`R3(config-router)#router-id 3.3.3.3`	Set the router ID for IP address `3.3.3.3`.
`R3(config-router)#transport-address ipv4 3.3.3.3`	Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter `ipv6` if you are configuring an IPv6 interface.
`R3(config-router)#neighbor all tcp-mss 720`	Configure the TCP MSS value with `all` neighbor.
`R3(config-router)#targeted-peer ipv4 1.1.1.1`	Configure targeted peer.
`R3(config-router-targeted-peer)#exit`	Exit-targeted-peer-mode.
`R3(config-router)#exit`	Exit the Router mode and return to the Configure mode.
`R3(config)#interface xe1`	Enter interface mode.
`R3(config-if)#enable-ldp ipv4`	Enable LDP on `xe1`.
`R3(config-if)#commit`	Commit the transaction.

R3 - OSPF

`R3(config)#router ospf 100`	Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process.
`R3(config-router)#network 20.20.20.0/24 area 0`	Define the interfaces `20.20.20.0/24`, on which OSPF runs and associate the area ID (0) with them.
`R3(config-router)#network 3.3.3.3/32 area 0`	Define the interfaces `3.3.3.3/32`, on which OSPF runs and associate the area ID (0) with them.
`R3(config-router)#commit`	Commit the transaction.

Validation

R1

```
R1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
       2.2.2.2            xe1        Passive   OPERATIONAL   30    00:11:22
       3.3.3.3            xe1        Passive   OPERATIONAL   30    00:11:22
R1#show ldp session 2.2.2.2
Session state                  : OPERATIONAL
Session role                   : Passive
TCP Connection                 : Established
IP Address for TCP             : 2.2.2.2
Interface being used           : xe1
Peer LDP ID                    : 2.2.2.2:0
Preferred Peer LDP Password    : Not Set
Adjacencies                    : 10.10.10.2
Advertisement mode             : Downstream Unsolicited
Label retention mode           : Liberal
Graceful Restart               : Not Capable
Keepalive Timeout              : 30
Reconnect Interval             : 15
Configured TCP MSS             : 700
Applied TCP MSS                : 700
Preferred TCP MSS              : NA
Address List received          : 2.2.2.2
                                 10.10.10.2
                                 20.20.20.1
Received Labels :       Fec              Label           Maps To
                 IPV4:20.20.20.0/24      impl-null       none
                 IPV4:10.10.10.0/24      impl-null       none
                 IPV4:3.3.3.3/32         25601           none
                 IPV4:2.2.2.2/32         impl-null       none
Sent Labels :    Fec                     Label           Maps To
                 IPV4:10.10.10.0/24      impl-null       none
                 IPV4:1.1.1.1/32         impl-null      none
R1#show ldp session 3.3.3.3
Session state                  : OPERATIONAL
Session role                   : Passive
TCP Connection                 : Established
IP Address for TCP             : 3.3.3.3
Interface being used           : xe1
Peer LDP ID                    : 3.3.3.3:0
Preferred Peer LDP Password    : Not Set
Adjacencies                    : 3.3.3.3
Advertisement mode             : Downstream Unsolicited
```

```
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
Configured TCP MSS            : 700
Applied TCP MSS               : 700
Preferred TCP MSS             : NA
Address List received         : 3.3.3.3
                                20.20.20.2
Received Labels :       Fec             Label           Maps To
Sent Labels :   Fec                     Label           Maps To
```

R2

```
R2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
       3.3.3.3            xe1        Passive   OPERATIONAL   30    00:13:39
       1.1.1.1            xe2       Active    OPERATIONAL   30    00:13:39
R2#show ldp session 3.3.3.3
Session state                 : OPERATIONAL
Session role                  : Passive
TCP Connection                : Established
IP Address for TCP            : 3.3.3.3
Interface being used          : xe1
Peer LDP ID                   : 3.3.3.3:0
Preferred Peer LDP Password   : Not Set
Adjacencies                   : 20.20.20.2
Advertisement mode            : Downstream Unsolicited
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
Configured TCP MSS            : 710
Applied TCP MSS               : 710
Preferred TCP MSS             : NA
Address List received         : 3.3.3.3
                                20.20.20.2
Received Labels :       Fec             Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
                IPV4:3.3.3.3/32         impl-null       25601
Sent Labels :   Fec                     Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
                IPV4:10.10.10.0/24      impl-null       none
                IPV4:2.2.2.2/32         impl-null       none
                IPV4:1.1.1.1/32         25600           impl-null
R2#show ldp session 1.1.1.1
Session state                 : OPERATIONAL
```

```
Session role                  : Active
TCP Connection                : Established
IP Address for TCP            : 1.1.1.1
Interface being used          : xe2
Peer LDP ID                   : 1.1.1.1:0
Preferred Peer LDP Password   : Not Set
Adjacencies                   : 10.10.10.1
Advertisement mode            : Downstream Unsolicited
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
Configured TCP MSS            : 710
Applied TCP MSS               : 700
Preferred TCP MSS             : NA
Address List received         : 1.1.1.1
                                10.10.10.1
Received Labels :       Fec             Label           Maps To
                IPV4:48.48.48.48/32     impl-null       none
                IPV4:10.10.10.0/24      impl-null       none
                IPV4:1.1.1.1/32         impl-null       25600
Sent Labels :   Fec                     Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
                IPV4:10.10.10.0/24      impl-null       none
                IPV4:3.3.3.3/32         25601           impl-null
                IPV4:2.2.2.2/32         impl-null       none
```

R3

```
R3#show ldp session 2.2.2.2
Session state                 : OPERATIONAL
Session role                  : Active
TCP Connection                : Established
IP Address for TCP            : 2.2.2.2
Interface being used          : xe1
Peer LDP ID                   : 2.2.2.2:0
Preferred Peer LDP Password   : Not Set
Adjacencies                   : 20.20.20.1
Advertisement mode            : Downstream Unsolicited
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
Configured TCP MSS            : 720
Applied TCP MSS               : 710
Preferred TCP MSS             : NA
Address List received         : 2.2.2.2
                                10.10.10.2
                                20.20.20.1
Received Labels :       Fec             Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
```

```
                IPV4:10.10.10.0/24      impl-null       none
                IPV4:2.2.2.2/32         impl-null       none
                IPV4:1.1.1.1/32         25600           none
Sent Labels :   Fec                     Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
                IPV4:3.3.3.3/32         impl-null       none
R3#show ldp session 1.1.1.1
Session state                 : OPERATIONAL
Session role                  : Active
TCP Connection                : Established
IP Address for TCP            : 1.1.1.1
Interface being used          : xe1
Peer LDP ID                   : 1.1.1.1:0
Preferred Peer LDP Password   : Not Set
Adjacencies                   : 1.1.1.1
Advertisement mode            : Downstream Unsolicited
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
Configured TCP MSS            : 720
Applied TCP MSS               : 700
Preferred TCP MSS             : NA
Address List received         : 1.1.1.1
                                10.10.10.1
Received Labels :       Fec             Label           Maps To
Sent Labels :   Fec                     Label           Maps To
```

Configuration of TCP MSS with Auto-targeted

R1 - NSM

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R1(config-if)#ip address 10.10.10.1/24`	Assign IP address `10.10.10.1/24` to interface.
`R1(config-if)#label-switching`	Enable label switching on interface `xe1`.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface lo`	Specify the loopback interface to be configured.
`R1(config-if)#ip address 1.1.1.1/32 secondary`	Set the IP address of the loopback interface to `1.1.1.1/32`.
`R1(config-if)#commit`	Commit the transaction.

R1 - LDP

`R1(config)#router ldp`	Enter Router mode for LDP.
`R1(config-router)#router-id 1.1.1.1`	Set the router ID to IP address `1.1.1.1`.

`R1(config-router)#transport-address ipv4 1.1.1.1`	Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter `ipv6` if you are configuring an IPv6 interface.
`R1(config-router)#targeted-peer ipv4 3.3.3.3`	Configure targeted peer.
`R1(config-router-targeted-peer)#exit`	Exit-targeted-peer-mode.
`R1(config-router)#exit`	Exit the Router mode and return to the configure mode.
`R1(config)#interface xe1`	Enter interface mode.
`R1(config-if)#enable-ldp ipv4`	Enable LDP on `xe1`.
`R1(config-if)#commit`	Commit the transaction.

R1 - OSPF

`R1(config)#router ospf 100`	Configure the routing process and specify the process ID (100). The process ID should be a unique positive integer identifying the routing process.
`R1(config-router)#network 10.10.10.0/24 area 0`	Define the interface `10.10.10.0/24`, on which OSPF runs and associate the area ID (0) with the interface.
`R1(config-router)#network 1.1.1.1/32 area 0`	Define the interface `1.1.1.1/32`, on which OSPF runs and associate the area ID (0) with the interface.
`R1(config-router)#commit`	Commit the transaction.

R2 - NSM

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface lo`	Specify the loopback interface to be configured.
`R2(config-if)#ip address 2.2.2.2/32 secondary`	Set the IP address of the loopback interface to `2.2.2.2/32`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R2(config-if)#ip address 20.20.20.1/24`	Assign IP address `20.20.20.1/24` to interface.
`R2(config-if)#label-switching`	Enable label switching on interface `xe1`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe2`	Specify the interface `xe2` to be configured.
`R2(config-if)#ip address 10.10.10.2/24`	Assign IP address `10.10.10.2/24` to interface.
`R2(config-if)#label-switching`	Enable label switching on interface `xe2`.
`R2(config-if)#commit`	Commit the transaction.

R2 - LDP

`R2(config)#router ldp`	Enter Router mode.
`R2(config-router)#router-id 2.2.2.2`	Set the router ID to IP address `2.2.2.2`.

`R2(config-router)#transport-address ipv4 2.2.2.2`	Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter `ipv6` if you are configuring an IPv6 interface.
`R2(config-router)#neighbor auto-targeted tcp-mss 800`	Configure the TCP MSS value on all `auto-targeted` neighbors.
`R2(config-router)#exit`	Exit Router mode and return to configure mode.
`R2(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R2(config-if)#enable-ldp ipv4`	Enable LDP on a specified interface `xe1`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe2`	Specify the interface `xe2` to be configured.
`R2(config-if)#enable-ldp ipv4`	Enable LDP on a specified interface `xe2`.
`R2(config-if)#commit`	Commit the transaction.

R2 - OSPF

`R2(config)#router ospf 100`	Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process.
`R2(config-router)#network 10.10.10.0/24 area 0`	Define the interfaces `10.10.10.0/24`, on which OSPF runs and associate the area ID (0) with them.
`R2(config-router)#network 20.20.20.0/24 area 0`	Define the interfaces `20.20.20.0/24`, on which OSPF runs and associate the area ID (0) with them.
`R2(config-router)#network 2.2.2.2/32 area 0`	Define the interfaces `2.2.2.2/32`, on which OSPF runs and associate the area ID (0) with them.
`R2(config-router)#commit`	Commit the transaction.

R3 - NSM

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface lo`	Specify the loopback interface to be configured.
`R3(config-if)#ip address 3.3.3.3/32 secondary`	Set the IP address of the loopback interface to `3.3.3.3/32`.
`R3(config-if)#exit`	Exit interface mode.
`R3(config)#interface xe1`	Specify the interface `xe1` to be configured.
`R3(config-if)#ip address 20.20.20.2/24`	Set the IP address of the interface to `20.20.20.2/24`.
`R3(config-if)#label-switching`	Enable label switching on interface `xe1`.
`R3(config-if)#commit`	Commit the transaction.

R3 - LDP

`R3(config)#router ldp`	Enter Router mode.
`R3(config-router)#router-id 3.3.3.3`	Set the router ID for IP address `3.3.3.3`.

`R3(config-router)#transport-address ipv4 3.3.3.3`	Configure the transport address to be used for a TCP session over which LDP will run on an IPv4 interface. Note: It is preferable to use the loopback address as transport address. In addition, use the parameter `ipv6` if you are configuring an IPv6 interface.
`R3(config-router)#neighbor auto-targeted tcp-mss 810`	Configure the TCP MSS value on all `auto-targeted` neighbors.
`R3(config-router)#targeted-peer ipv4 1.1.1.1`	Configure targeted peer.
`R3(config-router-targeted-peer)#exit`	Exit-targeted-peer-mode.
`R3(config-router)#exit`	Exit the Router mode and return to the configure mode.
`R3(config)#interface xe1`	Enter interface mode `xe1`.
`R3(config-if)#enable-ldp ipv4`	Enable LDP on `xe1`.
`R3(config-if)#commit`	Commit the transaction.

R3 - OSPF

`R3(config)#router ospf 100`	Configure the routing process and specify the Process ID (100). The Process ID should be a unique positive integer identifying the routing process.
`R3(config-router)#network 20.20.20.0/24 area 0`	Define the interfaces `20.20.20.0/24`, on which OSPF runs and associate the area ID (0) with them.
`R3(config-router)#network 3.3.3.3/32 area 0`	Define the interfaces `3.3.3.3/32`, on which OSPF runs and associate the area ID (0) with them.
`R3(config-router)#commit`	Commit the transaction.

Validation

R1

```
R1#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
Code  Peer IP Address      IF Name    My Role    State       KeepAlive  UpTime
       2.2.2.2            xe1        Passive   OPERATIONAL   30    00:00:03
       3.3.3.3            xe1        Passive   OPERATIONAL   30    00:00:03
R1#show ldp targeted-peers
IP Address          Interface
3.3.3.3             xe1
R1#show ldp session 3.3.3.3
Session state                 : OPERATIONAL
Session role                  : Passive
TCP Connection                : Established
IP Address for TCP            : 3.3.3.3
Interface being used          : xe1
```

```
Peer LDP ID                   : 3.3.3.3:0
Preferred Peer LDP Password   : Not Set
Adjacencies                   : 3.3.3.3
Advertisement mode            : Downstream Unsolicited
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
Configured TCP MSS            : Not configured
Applied TCP MSS               : 810
Preferred TCP MSS             : NA
Address List received         : 3.3.3.3
                                20.20.20.2
Received Labels :       Fec             Label           Maps To
                IPV4:20.20.20.0/24      25604           none
                IPV4:3.3.3.3/32         25603           none
                IPV4:10.10.10.0/24      25602           none
                IPV4:2.2.2.2/32         25601           none
                IPV4:1.1.1.1/32         25600           none
Sent Labels :   Fec                     Label           Maps To
                IPV4:10.10.10.0/24      25604           none
                IPV4:1.1.1.1/32         25603           none
                IPV4:20.20.20.0/24      25602           impl-null
                IPV4:3.3.3.3/32         25601           25601
                IPV4:2.2.2.2/32         25600           impl-null
```

R2

```
R2#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
       3.3.3.3            xe1       Passive   OPERATIONAL   30    00:00:04
       1.1.1.1            xe2       Active    OPERATIONAL   30    00:00:04
R2#show ldp targeted-peers
R2#show ldp session 3.3.3.3
Session state                 : OPERATIONAL
Session role                  : Passive
TCP Connection                : Established
IP Address for TCP            : 3.3.3.3
Interface being used          : xe1
Peer LDP ID                   : 3.3.3.3:0
Preferred Peer LDP Password   : Not Set
Adjacencies                   : 20.20.20.2
Advertisement mode            : Downstream Unsolicited
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
```

```
Configured TCP MSS            : Not configured
Applied TCP MSS               : 1460
Preferred TCP MSS             : NA
Address List received         : 3.3.3.3
                                20.20.20.2
Received Labels :       Fec             Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
                IPV4:3.3.3.3/32         impl-null       25601
Sent Labels :   Fec                     Label           Maps To
                IPV4:20.20.20.0/24      impl-null       none
                IPV4:10.10.10.0/24      impl-null       none
                IPV4:2.2.2.2/32         impl-null       none
                IPV4:1.1.1.1/32         25600           impl-null
```

R3

```
R3#show ldp session
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually
Code  Peer IP Address      IF Name    My Role    State      KeepAlive  UpTime
       2.2.2.2            xe1        Active    OPERATIONAL   30    00:02:15
       1.1.1.1            xe1        Active    OPERATIONAL   30    00:02:15
R3#show ldp targeted-peers
IP Address          Interface
1.1.1.1             xe1
PE2#show ldp session 1.1.1.1
Session state                 : OPERATIONAL
Session role                  : Active
TCP Connection                : Established
IP Address for TCP            : 1.1.1.1
Interface being used          : xe1
Peer LDP ID                   : 1.1.1.1:0
Preferred Peer LDP Password   : Not Set
Adjacencies                   : 1.1.1.1
Advertisement mode            : Downstream Unsolicited
Label retention mode          : Liberal
Graceful Restart              : Not Capable
Keepalive Timeout             : 30
Reconnect Interval            : 15
Configured TCP MSS            : 810
Applied TCP MSS               : 810
Preferred TCP MSS             : NA
Address List received         : 1.1.1.1
                                10.10.10.1
Received Labels :       Fec             Label           Maps To
                                IPV4:10.10.10.0/24      25604           none
                IPV4:1.1.1.1/32         25603           none
                IPV4:20.20.20.0/24      25602           none
                IPV4:3.3.3.3/32         25601           none
```

```
                IPV4:2.2.2.2/32         25600           none
Sent Labels :   Fec                     Label           Maps To
                IPV4:20.20.20.0/24      25604           none
                IPV4:3.3.3.3/32         25603           none
                IPV4:10.10.10.0/24      25602           impl-null
                IPV4:2.2.2.2/32         25601           impl-null
                IPV4:1.1.1.1/32         25600           25600
```

New CLI Command

neighbor tcp-mss

Use this command to set the TCP MSS for an LDP session. MSS is a TCP parameter that defines the maximum amount of data in a TCP segment that can be transmitted.

Use the `no` command to remove the TCP MSS from an LDP session.

Command Syntax

```
neighbor (A.B.C.D| auto-targeted | all) tcp-mss <560-1440>
no neighbor (A.B.C.D | auto-targeted | all) tcp-mss
```

Parameters

`A.B.C.D`	To set MSS for the specific peer.
`auto-targeted`	To set MSS for auto-targeted LDP peer. Auto-targeted LDP sessions automatically establish the TCP connection with neighboring routers and do not require the manual configuration of each peer.
`all`	To set MSS for all LDP peers
`<560-1440>`	Configure the TCP MSS between this range.

Default

By default, `neighbor tcp-mss` is disabled and the MSS value is 1460 bytes.

Command Mode

Router LDP mode.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
OcNOS(config)#router ldp
OcNOS(config-router)#neighbor 2.2.2.2 tcp-mss 900
OcNOS(config-router)#neighbor all tcp-mss 1000
OcNOS(config-router)#neighbor auto-targeted tcp-mss 800
OcNOS(config-router)#commit
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
ACK	Acknowledgment
IGP	Interior Gateway Protocol
LDP	Label Distribution Protocol
MSS	Maximum Segment Size
MTU	Maximum Transmission Unit
OSPF	Open Short Path First
SYN	Synchronize
TCP	Transmission Control Protocol

Glossary

The following provides definitions for key terms used throughout this document:

LDP	LDP is a routing protocol that manages and distributes the labels to the route in a Multiprotocol Label Switching (MPLS) network. Adding a label to a route helps to control the flow of network traffic and increases the forwarding speed, ensuring a smooth and optimized data transmission.
LDP session	LDP session is the connection established between LDP routers in an MPLS network.
MSS	MSS is a TCP parameter that defines the maximum amount of data in a TCP segment that can be transmitted.
TCP	TCP is one of the main protocols in the Internet Protocol (IP) suite. It offers a secure and reliable connection between two devices.
TCP segment	TCP segment is a unit of data transmitted in a TCP connection. The segment consists of header and payload. The header contains the control information to manage the transmission, and the payload contains the actual data that needs to be transmitted.

Single Home VxLAN IRB with OSPF or ISIS

Overview

Single Home Virtual Extensible LAN (VxLAN) with Integrated Routing (IRB) using Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (ISIS) protocols provides the solution for connecting and managing virtual networks within a data center or network infrastructure.

This feature offers a solution for networks where the interconnection of VLANs is required. These protocols can be configured on IRB interfaces within layer 3 switches or routers. This configuration enables dynamic routing, facilitating the exchange of routing information with other devices in the network. By assigning IP addresses to the IRB interfaces, they serve as the default gateways for devices within the respective VLANs.

Both OSPF and ISIS routing updates are dynamically exchanged over IRB interfaces, ensuring up-to-date routing tables and optimized traffic routing across different VLANs and networks.

This feature offers flexibility in configuring network topologies, and ensures compatibility and interoperability within diverse network environments.

Feature Characteristics

The OSPF and ISIS support over the IRB Interface feature has the following characteristics:

- Enables the control of Receive (RX)/ Transmit (TX) of OSPF and ISIS packets on IRB interfaces, providing effective management of IRB interfaces interactions with OSPF and ISIS for optimized network communication and routing.
- IRB interfaces process configured MTU size packets.
- Maintains consistency in CLI commands with SVI interfaces for OSPF and ISIS configurations, simplifying network management tasks.

Benefits

The OSPF and ISIS support over the IRB Interface has the following benefits:

- Enables seamless inter-subnet communication across different VNIDs and subnets within the same customer network.
- Promotes seamless connectivity between devices, irrespective of whether they are connected through IRB or SVI interfaces, and simplifies network management.
- The network gains greater adaptability to various scenarios and evolving requirements, offering greater versatility in its operations.

Prerequisites

- Router must be up and running.
- Maintain synchronization with VRF changes by performing IRB shut/no shut actions when specific events occur within the IPVRF. These events may involve adding or removing Route Targets (RTs), updating Route Distinguishers (RDs), or modifying Layer 3 Virtual Network Identifiers (L3VNIs).

Topology for OSPF

The network topology includes various network elements such as routers, customer edge (CE) devices, Service Aggregator (SA) devices, and Provider Edge (PE) routers. The feature enables OSPF on the IRB interfaces, allowing for efficient routing and communication between network devices within the topology.

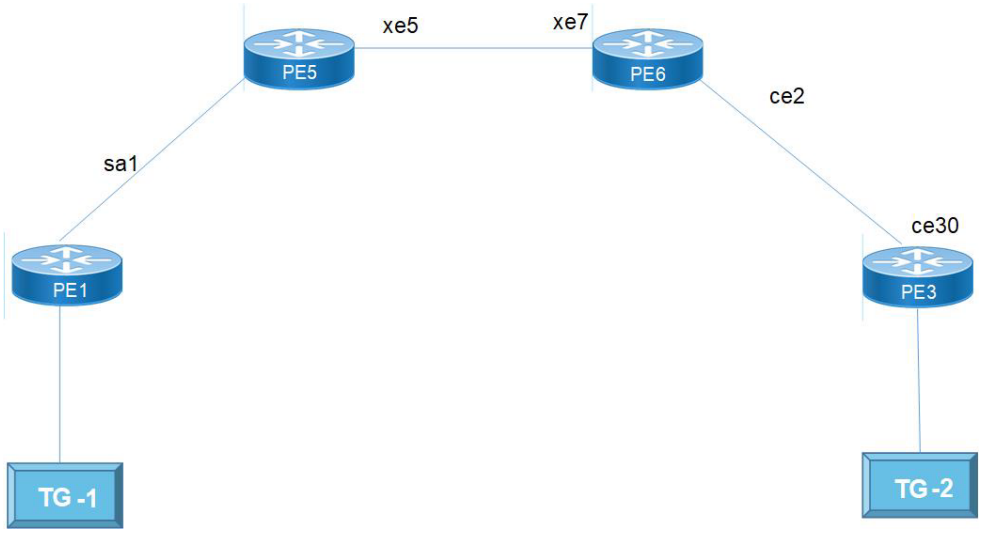

Single Home VxLAN IRB with OSPF

Configuration

Perform the following configurations to set up different interfaces, routing protocols, and BGP parameters to enable VXLAN, IRB, and EVPN functionality in the network.

Configure OSPF

PE1

`PEI(Config)# terminal`	Enters the configuration mode.
`PE1(config)#interface sa1`	Configure the sa1 interface as a network interface.
`PE1(config-if)# ip address 10.1.1.1/24`	Assigns an IP address to the sa1 interface with a subnet mask of /24.
`PE1(config-if)# ip ospf cost 10`	Configures the OSPF cost for the sa1 interface, setting it to 10.
`PE1(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa1 interface.
`PE1(config)#interface xe1`	Enters the interface xe1 mode.
`PE1(config-if)# static-channel-group 1`	Assigns the static channel group 1 to the xe1 interface.
`PE1(config-irb-if)#interface lo`	Configures the loopback (lo) interface.
`PE1(config-if)# ip address 1.1.1.1/32 secondary`	Assigns the primary IP address 1.1.1.1/32 to the loopback interface and specifies it as secondary.

`PE1(config)#router ospf 1`	Enters the OSPF configuration mode for OSPF process 1.
`PE1(config-router)# ospf router-id 1.1.1.1`	Sets the OSPF router ID to 1.1.1.1 for OSPF process 1.
`PE1(config-router)# network 1.1.1.1/32 area 0.0.0.0`	Advertises the network 1.1.1.1/32 into OSPF area 0.0.0.0.
`PE1(config-router)# network 10.1.1.0/24 area 0.0.0.0`	Advertises the network 10.1.1.0/24 into OSPF area 0.0.0.0.
`PE1(config)#nvo vxlan enable`	Enables the VXLAN feature on the device.
`PE1(config)#nvo vxlan irb`	Enables VXLAN IRB functionality.
`PE1(config-vrf)#mac vrf L2VRF1`	Configures a MAC VRF named L2VRF1.
`PE1(config-vrf)# rd 1.1.1.1:11`	Sets the Route Distinguisher (RD) to 1.1.1.1:11 for the VRF.
`PE1(config-vrf)# route-target both 9.9.9.9:100`	Configures both import and export route targets for the VRF.
`PE1(config-vrf)#ip vrf L3VRF1`	Configures an IP VRF named L3VRF1.
`PE1(config-vrf)# rd 51000:11`	Sets the RD value to 51000:11 for the L3VRF1.
`PE1(config-vrf)# route-target both 100:100`	Configures both import and export route targets for L3VRF1.
`PE1(config-vrf)# l3vni 1000`	Configures the L3 Virtual Network Identifier (L3VNI) with the value 1000.
`PE1(config)#interface irb1001`	Configures the IRB interface for L3VRF1.
`PE1(config-irb-if)# ip vrf forwarding L3VRF1`	Assigns the L3VRF1 to the IRB interface.
`PE1(config-irb-if)# ip address 11.11.11.1/24`	Assigns an IP address 11.11.11.1/24 to the IRB interface.
`PE1(config-irb)#interface irb2001`	Configures the IRB interface for IPv6 in L3VRF1.
`PE1(config-irb-if)# ip vrf forwarding L3VRF1`	Assigns the L3VRF1 to the IPv6 IRB interface.
`PE1(config-irb-if)# ipv6 address 2001::1/64`	Assigns an IP address 11.11.11.1/24 to the IRB interface.
`PE1(config-irb-if)#mtu 9000`	Sets the Maximum Transmission Unit (MTU) for this IRB interface to 9000 bytes.
`PE1(config-router)#router ospf 2 L3VRF1`	Configures OSPF on the L3VRF1.
`PE1(config-router)# network 11.11.11.0/24 area 0.0.0.0`	Advertises the network 11.11.11.0/24 into OSPF area 0.0.0.0.
`PE1(config-router)#router ipv6 vrf ospf L3VRF1`	Configures OSPFv3 on the L3VRF1.
`PE1(config-router)# router-id 1.1.1.1`	Configures the router ID as 1.1.1.1.
`PE1(config-irb)#interface irb2001`	Configures the IPv6 IRB interface.
`PE1(config-irb-if)# ipv6 router ospf area 0.0.0.0 tag L3VRF1 instance-id 0`	Attaches the OSPFv3 instance ID to the IPv6 IRB interface.
`PE1(config)#nvo vxlan vtep-ip-global 1.1.1.1`	Configures the global VTEP IP address as 1.1.1.1.
`PE1(config)#nvo vxlan id 101 ingress-replication`	Configures the VXLAN ID as 101 for ingress replication.
`PE1(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Maps the EVPN-BGP host reachability protocol to L2VRF1.
`PE1(config-nvo)# evpn irb1001`	Maps the IRB interface 1001 to EVPN.
`PE1(config-nvo)# vni-name VNI-101`	Configures the VNI name as VNI-101.
`PE1(config)#nvo vxlan id 2001 ingress-replication`	Configures the VXLAN ID as 2001 for ingress replication.
`PE1(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Maps the EVPN-BGP host reachability protocol to L2VRF1.

`PE1(config-nvo)# evpn irb2001`	Maps the IPv6 IRB interface to EVPN.
`PE1(config)#interface xe2`	Configures the xe2 interface.
`PE1(config-if)# switchport`	Configures the port as a Layer 2 (L2) switchport.
`PE1(config-if)# load-interval 30`	Configures the load-interval of 30 minutes for monitoring traffic on the xe2 interface.
`PE1(config)#nvo vxlan access-if port-vlan xe2 100`	Configures a VxLAN network virtualization overlay (NVO) on the interface xe2 with VLAN ID 100
`PE1(config-nvo-acc-if)# map vnid 101`	Maps VLAN 100 to the VxLAN Network Identifier (VNID) 101.
`PE1(config-nvo-acc-if)#nvo vxlan access-if port-vlan xe2 2001`	Configures another VxLAN NVO on the same interface xe2, but this time with VLAN ID 2001
`PE1(config-nvo-acc-if)# map vnid 2001`	Maps VLAN 2001 to a different VxLAN VNID.
`PE1(config-router)#router bgp 100`	Configures the BGP process with AS number 100.
`PE1(config-router)# bgp router-id 1.1.1.1`	Assigns the router ID as 1.1.1.1 for the BGP instance.
`PE1(config-router)# neighbor 4.4.4.4 remote-as 100`	Configures neighbor 4.4.4.4 with a remote AS number of 100.
`PE1(config-router)# neighbor 4.4.4.4 update-source lo`	Configures the update source for neighbor 4.4.4.4 to be the loopback interface.
`PE1(config-router)# neighbor 4.4.4.4 advertisement-interval 0`	Configures the advertisement interval for neighbor 4.4.4.4 as 0.
`PE1(config-router)# address-family l2vpn evpn`	Configures the address-family for L2VPN EVPN.
`PE1(config-router-af)# neighbor 4.4.4.4 activate`	Activates the neighbor for the L2VPN EVPN address-family.
`PE1(config-router-af)# exit-address-family`	Exits from the address family configuration.
`PE1(config-router)# address-family ipv4 vrf L3VRF1`	Configures the IPv4 address-family for VRF L3VRF1.
`PE1(config-router-af)# redistribute connected`	Configures the redistribution of connected routes within the IPv4 address-family.
`PE1(config-router-af)# exit-address-family`	Exits the IPv4 address-family configuration.
`PE1(config-router)# address-family ipv6 vrf L3VRF1`	Configures the IPv6 address-family for VRF L3VRF1.
`PE1(config-router-af)# redistribute connected`	Configures the redistribution of connected routes within the IPv6 address-family.
`PE1(config-router-af)# exit-address-family`	Exits the IPv6 address-family configuration.

PE5

`PE5#configure terminal`	Enters the configuration mode
`PE5(config)#interface sa1`	Configure the sa1 interface as a network interface.
`PE5(config-if)# ip address 10.1.1.1/24`	Assigns an IP address to the sa1 interface with a subnet mask of /24.
`PE5(config-if)# ip ospf cost 10`	Configures the OSPF cost for the sa1 interface, setting it to 10.
`PE5(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa1 interface.
`PE5(config)#interface xe1`	Configure network interface towards PE6.
`PE5(config-if)# static-channel-group 1`	Assigns the static channel group 1 to the xe1 interface.

`PE5(config)#interface xe5`	configures the xe5 interface.
`PE5(config-if)#ip address 30.1.1.1/24`	Assigns the primary IP address 1.1.1.1/32 to the loopback interface and specifies it as secondary.
`PE5(config)#ip ospf cost 10`	Configures the OSPF cost for the xe5 interface, setting it to 10.
`PE5(config-router)# ospf router-id 1.1.1.1`	Assigns an IP address (30.1.1.1) to the xe5 interface with a subnet mask of /24.
`PE5(config)#load-interval 30`	Configures the load-interval for monitoring traffic on the xe5 interface.
`PE5(config)#router ospf 1`	Enters the OSPF configuration mode for OSPF process 1.
`PE5(config-router)# network 30.1.1.0/24 area 0.0.0.0`	Advertises the network 30.1.1.0/24 into OSPF area 0.0.0.0.
`PE5(config-router)# network 10.1.1.0/24 area 0.0.0.0`	Advertises the network 10.1.1.0/24 into OSPF area 0.0.0.0.

PE3

`PE3#configure terminal`	Enters the configuration mode
`PE3(config)#interface  ce30`	Configure the ce30 interface as a network interface.
`PE3(config-if)# ip address 40.1.1.2/24`	Assigns an IP address to the ce30 interface with a subnet mask of /24.
`PE3(config-if)# ip ospf cost 10`	Configures the OSPF cost for the sa1 interface, setting it to 10.
`PE3(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa1 interface.
`PE3(config)#interface lo`	Configure the loopback interface.
`PE3(config-if)#ip address 4.4.4.4/32 secondary`	Assign an secondary IP to an loopback interface.
`PE3(config)#ip ospf cost 10`	Configures the OSPF cost for the xe7interface, setting it to 10.
`PE3(config)#load-interval 30`	Configures the load-interval for monitoring traffic on the xe5 interface.
`PE3(config)#router ospf 1`	Enters the OSPF configuration mode for OSPF process 1.
`PE3(config-router)# ospf router-id 4.4.4.4`	Configures the router id to an ospf instance.
`PE3(config-router)# network 4.4.4.4/32 area 0.0.0.0`	Advertises the loopback address.
`PE3(config-router)# network 40.1.1.0/24 area 0.0.0.0`	Advertises the network interface IP address.
`PE3(config)#nvo vxlan enable`	Enables VXLAN on the device, allowing it to participate in VXLAN networks.
`PE3(config)#nvo vxlan irb`	Enables VXLAN IRB functionality, that allows routing between VXLAN and non-VXLAN networks.
`PE3(config-vrf)#mac vrf L2VRF1`	Configures a L2 MAC VRF instance named L2VRF1, which is a logical network segment for L2 traffic isolation.
`PE3(config-vrf)# rd 4.4.4.4:11`	Configures a RD for the L2VRF1, with the value 4.4.4.4:11.
`PE3(config-vrf)# route-target both 9.9.9.9:100`	Configures a route target for the VRF.

`PE3(config-vrf)#ip vrf L3VRF1`	Configures a L3 VRF named L3VRF1.
`PE3(config-vrf)# rd 56000:11`	Configures a RD for the L3VRF1, with the value 56000:11.
`PE3(config-vrf)# route-target both 100:100`	Configures a route target for the VRF.
`PE3(config-vrf)# l3vni 1000`	Configures a L3VNI with the ID 1000 for the VRF.
`PE3(config)#interface irb1001`	Configures the IRB interface with the ID 1001.
`PE3(config-irb-if)# ip vrf forwarding L3VRF1`	Associates the IRB interface with the L3VRF1, ensuring that traffic from this interface is isolated within that VRF.
`PE3(config-irb-if)# ip address 12.12.12.1/24`	Assigns an IP address 12.12.12.1 with a subnet mask of /24 to the IRB interface, enabling it for L3 routing.
`PE3(config-irb-if)# mtu 1500`	Configures the MTU for the interface irb1001 to 1500 bytes.
`PE3(config)#interface irb2001`	Configures another IRB interface with the ID 2001.
`PE3(config-irb-if)# ip vrf forwarding L3VRF1`	Associates the IRB interface with the L3VRF1.
`PE3(config-irb-if)# ipv6 address 2002::1/64`	Assigns an IPv6 address 2002::1 with a subnet mask of /64 to the IRB interface, enabling it for IPv6 routing.
`PE3(config-irb-if)# mtu 1500`	Configures the MTU for the interface irb2001 to 1500 bytes.
`PE3(config-router)#router ospf 2 L3VRF1`	Configures the OSPF routing process on OSPF instance 2 for the L3VRF1.
`PE3(config-router)# network 12.12.12.0/24 area 0.0.0.0`	Advertises the network 12.12.12.0/24 to OSPF area 0.0.0.0.
`PE3(config-router)#router ipv6 vrf ospf L3VRF1`	Configures the OSPFv3 routing process on OSPFv3 instance for the L3VRF1.
`PE3(config-router)# router-id 4.4.4.4`	Sets the router ID for the OSPF/OSPFv3 instances to 4.4.4.4.
`PE3(config)#nvo vxlan vtep-ip-global 4.4.4.4`	Configures the global VTEP IP address as 4.4.4.4 for VXLAN.
`PE3(config)#nvo vxlan id 102 ingress-replication`	Configures the VXLAN with VNI ID 102 for ingress replication.
`PE3(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Maps the VXLAN configuration with the EVPN-BGP protocol and associates it with the L2VRF1.
`PE3(config-nvo)# evpn irb1001`	Maps the IRB interface irb1001 to the VXLAN.
`PE3(config-nvo)# vni-name VNI-101`	Configures the VNI name as VNI-101.
`PE3(config)#nvo vxlan id 2002 ingress-replication`	Configures another VXLAN with VNI ID 2002 for ingress replication.
`PE3(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF1`	Maps the VXLAN configuration with the EVPN-BGP protocol and associates it with the L2VRF1.
`PE3(config-nvo)# evpn irb2001`	Maps the IPv6 IRB interface irb2001 to the VXLAN.
`PE3(config)#interface sa4`	Configures interface sa4.
`PE3(config-if)# switchport`	Configures the interface as a switchport.
`PE3(config-if)# load-interval 30`	Sets the load interval for the interface to 30 seconds.
`PE3(config-if)# mtu 1500`	Configures the MTU for the interface to 1500 bytes.
`PE3(config)#interface xe1`	Configures interface xe1.
`PE3(config-if)# static-channel-group 4`	Assigns a static channel group to interface xe1.
`PE3(config)#nvo vxlan access-if port-vlan sa4 100`	Configures a VxLAN nNVO on the interface xe2 with VLAN ID 100

`PE3(config-nvo-acc-if)# map vnid 101`	Maps VLAN 100 to the VxLAN VNID 101.
`PE3(config-nvo-acc-if)#nvo vxlan access-if port-vlan sa4 2001`	Configures another VxLAN NVO on the same interface xe2.
`PE3(config-nvo-acc-if)# map vnid 2001`	Maps VLAN 2001 to a different VxLAN VNID, in this case, VNID 2001.
`PE3(config-router)#router bgp 100`	Configures the BGP with AS number 100.
`PE3(config-router)# bgp router-id 4.4.4.4`	Sets the BGP router ID to 4.4.4.4.
`PE3(config-router)# neighbor 1.1.1.1 remote-as 100`	Configures a BGP neighbor with the remote AS number 100 and the IP address 1.1.1.1.
`PE3(config-router)# neighbor 1.1.1.1 update-source lo`	Specifies the BGP neighbor to use the loopback interface as the source for updates.
`PE3(config-router)# neighbor 1.1.1.1 advertisement-interval 0`	Configures the advertisement interval for BGP neighbor updates.
`PE3(config-router)# address-family l2vpn evpn`	Configures the BGP address family for Layer 2 VPN EVPN.
`PE3(config-router-af)# neighbor 1.1.1.1 activate`	Activates the BGP neighbor for the specified address family.
`PE3(config-router-af)# exit-address-family`	Exits the BGP address family configuration.
`PE3(config-router)# address-family ipv4 vrf L3VRF1`	Configures the BGP address family for IPv4 within VRF L3VRF1.
`PE3(config-router-af)# redistribute connected`	Configures BGP to redistribute connected routes into the BGP process.
`PE3(config-router-af)# exit-address-family`	Exits the BGP address family configuration for IPv4.
`PE3(config-router)# address-family ipv6 vrf L3VRF1`	Configures the BGP address family for IPv6 within VRF L3VRF1.
`PE3(config-router-af)# redistribute connected`	Configures BGP to redistribute connected routes into the BGP process.
`PE3(config-router-af)# exit-address-family`	Exits the BGP address family configuration for IPv6.

PE6

`PE6#configure terminal`	Enters the configuration mode.
`PE6(config)#interface  ce2`	Configure the ce2 interface as a network interface.
`PE6(config-if)# ip address 10.1.1.1/24`	Assigns an IP address to the sa1 interface with a subnet mask of /24.
`PE6(config-if)# ip ospf cost 10`	Configures the OSPF cost for the sa1 interface, setting it to 10.
`PE6(config-if)# load-interval 30`	Configures the load-interval for monitoring traffic on the sa1 interface.
`PE6(config)#interface xe7`	Configure network interface towards PE5.
`PE6(config-if)# static-channel-group 1`	Assigns the static channel group 1 to the xe1 interface.
`PE6(config-if)#ip address 30.1.1.1/24`	Assign IP address to network interface.
`PE6(config)#ip ospf cost 10`	Configures the OSPF cost for the xe7interface, setting it to 10.

`PE6(config)#load-interval 30`	Configures the load-interval for monitoring traffic on the xe5 interface.
`PE6(config)#router ospf 1`	Enters the OSPF configuration mode for OSPF process 1.
`PE6(config-router)# network 30.1.1.0/24 area 0.0.0.0`	Advertises the network 30.1.1.0/24 into OSPF area 0.0.0.0.
`PE6(config-router)# network 40.1.1.0/24 area 0.0.0.0`	Advertises the network 40.1.1.0/24 into OSPF area 0.0.0.0.

Topology for ISIS

The network topology includes various network elements such as routers, customer edge (CE) devices, Service Aggregator (SA) devices, and Provider Edge (PE) routers. The feature enables OSPF and ISIS support on the IRB interfaces, allowing for efficient routing and communication between network devices within the topology.

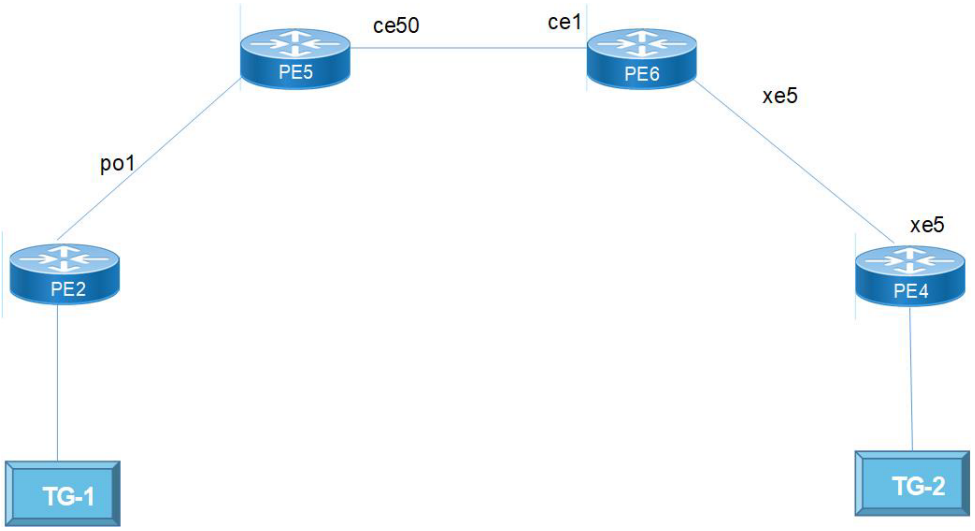

Single Home VxLAN IRB with ISIS

Configure ISIS

PE2

`PE2(config-if)# interface po1`	Enters configuration mode for po 1.
`PE2(config-if)# ip address 20.1.1.1/24`	Assigns the IP address 20.1.1.1 with a subnet mask of 255.255.255.0 to the interface.
`PE2(config-if)#ip router isis 1`	Enables ISIS routing protocol on the interface with process ID 1.
`PE2(config-if)#load-interval 30`	Sets the interval for which interface statistics are collected to 30 seconds.
`PE2(config)#nvo vxlan enable`	Enables the VXLAN feature on the device.
`PE2(config)#nvo vxlan irb`	Enables VXLAN IRB functionality.
`PE2(config-vrf)#mac vrf L2VRF2`	Enters the configuration mode for a MAC VRF named L2VRF2.

`PE2(config-vrf)# rd 2.2.2.2:11`	Sets the route distinguisher (RD) for the VRF to 2.2.2.2:11.
`PE2(config-vrf)#route-target both 10.10.10.10:100`	Specifies import and export route targets for the VRF.
`PE2(config-vrf)#ip vrf L3VRF2`	Enters the configuration mode for an IP VRF named L3VRF2.
`PE2(config-vrf)#rd 61000:11`	Sets the RD for the IP VRF to 61000:11
`PE2(config-vrf)# route-target both 101:101`	Specifies import and export route targets for the IP VRF.
`PE2(config-vrf)# l3vni 2000`	Configures the Layer 3 VNI (Virtual Network Identifier) for the IP VRF.
`PE2(config)#interface irb2001`	Enters the configuration mode for interface IRB2001.
`PE2(config-irb-if)# ip vrf forwarding L3VRF2`	Associates the interface with the IP VRF L3VRF2.
`PE2(config-irb-if)# ip address 13.13.13.1/24`	Configures an IP address with a subnet mask of /24 on IRB2001.
`PE2(config-irb-if)#mtu 9000`	Sets the Maximum Transmission Unit (MTU) for the interface to 9000 bytes.
`PE2(config-irb-if)#ip router isis 2`	Associates the interface with ISIS routing process 2.
`PE2(config-irb)#interface irb3001`	Enters the configuration mode for interface IRB3001.
`PE2(config-irb-if)# ip vrf forwarding L3VRF2`	Associates the interface with the IP VRF L3VRF2.
`PE2(config-irb-if)# ipv6 address 3001::1/64`	Configures an IPv6 address on IRB3001 with the specified prefix length.
`PE2(config-irb-if)#mtu 9000`	Sets the MTU for the interface to 9000 bytes.
`PE2(config-irb)#ipv6 router isis 3`	Associates the interface with IPv6 ISIS routing process 3.
`PE2(config)#router isis 2 L3VRF2`	Enters the configuration mode for ISIS routing process 2 within VRF L3VRF2.
`PE2(config-router)#is-type level-1-2`	Specifies the ISIS level type as level-1-2.
`PE2(config-router)#metric-style wide`	Configures a wide metric style for ISIS.
`PE2(config-router)# dynamic-hostname`	Enables dynamic hostname assignment for the ISIS router.
`PE2(config-router)# bfd all-interfaces`	Enables Bidirectional Forwarding Detection (BFD) on all interfaces within ISIS.
`PE2(config-router)#net 49.0000.0000.0221.00`	Configures the network entity title (NET) for the ISIS process.
`PE2(config)#router isis 3 L3VRF2`	Enters the configuration mode for ISIS routing process 3 within VRF L3VRF2.
`PE2(config-router)#is-type level-1-2`	Specifies the ISIS level type as level-1-2.
`PE2(config-router)# metric-style wide`	Configures a wide metric style for ISIS.
`PE2(config-router)# dynamic-hostname`	Enables dynamic hostname assignment for the ISIS router.
`PE2(config-router)#bfd all-interfaces`	Enables BFD on all interfaces within ISIS.
`PE2(config-router)# net 49.0000.0000.0222.00`	Configures the network entity title (NET) for ISIS routing with the specified value.
`PE2(config)#nvo vxlan vtep-ip-global 2.2.2.2`	Configures the global VxLAN VTEP IP address to 2.2.2.2.
`PE2(config)#nvo vxlan id 201 ingress-replication`	Configures a VxLAN with VNI 201 and specifies ingress-replication for multicast traffic handling.
`PE2(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Specifies the EVPN-BGP host-reachability-protocol for the VxLAN with the VRF L2VRF2
`PE2(config-nvo)# evpn irb2001`	Enables EVPN IRB (Integrated Routing and Bridging) for VxLAN interface IRB2001.
`PE2(config-nvo)# vni-name VNI-201`	Assigns a name VNI-201 to the VxLAN VNI 201.

`PE2(config)#nvo vxlan id 3001 ingress-replication`	Configures another VxLAN with VNI 3001 and specifies ingress-replicationr for multicast traffic handling.
`PE2(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Specifies the EVPN-BGP host-reachability-protocol for the VxLAN with the VRF L2VRF2.
`PE2(config-nvo)# evpn irb3001`	Enables EVPN IRB for VxLAN interface IRB3001.
`PE2(config-if)#interface xe11`	Enters the configuration mode for the interface 11.
`PE2(config-if)#switchport`	Configures the interface as a Layer 2 switchport.
`PE2(config-if)#load-interval 30`	Sets the interval for which interface statistics are collected to 30 seconds.
`PE2(config)#nvo vxlan access-if port-vlan xe11 100`	Configures a VxLAN network virtualization overlay (NVO) on the interface xe2 with VLAN ID 100
`PE2(config-nvo-acc-if)# map vnid 101`	Maps VLAN 100 to the VxLAN Network Identifier (VNID) 101.
`PE2(config-nvo-acc-if)#nvo vxlan access-if port-vlan xe11 2001`	Configures another VxLAN NVO on the same interface xe2, but this time with VLAN ID 2001
`PE2(config-nvo-acc-if)# map vnid 101`	Maps VLAN 100 to the VxLAN Network Identifier (VNID) 101.
`PE2(config-nvo-acc-if)#nvo vxlan access-if port-vlan xe11 2001`	Configures another VxLAN NVO on the same interface xe2, but this time with VLAN ID 2001
`PE2(config-nvo-acc-if)# map vnid 2001`	Maps VLAN 2001 to a different VxLAN VNID.
`PE2(config-if)#router isis 1`	Starts the ISIS routing process with process ID 1.
`PE2(config-if)#is-type level-1-2`	Specifies that the router participates in both Level 1 and Level 2 routing.
`PE2(config-if)#metric-style wide`	Configures the metric style to be wide, enabling more flexibility in metric calculations.
`PE2(config-if)#mpls traffic-eng router-id 2.2.2.2`	Sets the MPLS Traffic Engineering router ID to 2.2.2.2.
`PE2(config-if)#mpls traffic-eng level-1`	Enables MPLS Traffic Engineering for Level 1 ISIS.
`PE2(config-if)#mpls traffic-eng level-2`	Enables MPLS Traffic Engineering for Level 2 ISIS.
`PE2(config-if)#dynamic-hostname`	Enables the dynamic hostname feature for ISIS.
`PE2(config-if)#bfd all-interfaces`	Configures Bidirectional Forwarding Detection on all interfaces.
`PE2(config-if)#net 49.0000.0000.0001.00`	Specifies the network entity title (NET) for ISIS.

BGP Configuration

`PE2(config)#router bgp 100`	Starts the BGP routing process with an autonomous system number (AS) of 100.
`PE2(config-router)#bgp router-id 2.2.2.2`	Sets the BGP router ID to 2.2.2.2.
`PE2(config-router)#neighbor 3.3.3.3 remote-as 100`	Configures a BGP neighbor with the IP address 3.3.3.3 and specifies the remote AS number as 100.
`PE2(config-router)#neighbor 3.3.3.3 update-source lo`	Specifies that loopback interface (lo) is the source for BGP updates to the neighbor.
`PE2(config-router)#neighbor 3.3.3.3 advertisement-interval 0`	Sets the advertisement interval to 0, which means updates will be sent immediately.
`PE2(config-router)#address-family ipv4 unicast`	Enters the configuration mode for the IPv4 unicast address family within the router configuration.
`PE2(config-router-af)#network 2.2.2.2/32`	Specifies that network 2.2.2.2 with a /32 subnet mask is part of the IPv4 unicast address family.

`PE2(config-router-af)#neighbor 3.3.3.3 activate`	Activates the neighbor with the IP address 3.3.3.3 for the IPv4 unicast address family.
`PE2(config-router-af)#exit-address-family`	Exits the configuration mode for the IPv4 unicast address family.
`PE2(config-router)#address-family l2vpn evpn`	Enters the configuration mode for the L2VPN EVPN address family within the router configuration.
`PE2(config-router-af)#neighbor 3.3.3.3 activate`	Activates the neighbor with the IP address 3.3.3.3 for the L2VPN EVPN address family.
`PE2(config-router-af)#exit-address-family`	Exits the configuration mode for the L2VPN EVPN address family.
`PE2(config-router)#address-family ipv4 vrf L3VRF2`	Enters the configuration mode for the IPv4 address family within the VRF named L3VRF2.
`PE2(config-router-af)#redistribute connected`	Configures the redistribution of directly connected routes into the IPv4 address family for the specified VRF.
`PE2(config-router-af)#exit-address-family`	Exits the configuration mode for the IPv4 address family within the VRF L3VRF2.
`PE2(config-router-af)#address-family ipv6 vrf L3VRF2`	Enters the configuration mode for the IPv6 address family within the VRF named L3VRF2.
`PE2(config-router-af)#redistribute connected`	Configures the redistribution of directly connected routes into the IPv6 address family for the specified VRF.
`PE2(config-router-af)#exit-address-family`	Exits the configuration mode for the IPv6 address family within the VRF L3VRF2.

PE5

`PE5(config-if)#interface po1`	Enters the configuration mode for po1.
`PE5(config-if)#ip address 20.1.1.2/24`	Assigns the IP address 20.1.1.2 with a subnet mask of /24 to this interface.
`PE5(config-if)#ip router isis 1`	Specifies that ISIS routing process 1 is enabled on this interface.
`PE5(config-if)#load-interval 30`	Sets the load interval to 30 seconds for monitoring the interface.
`PE5(config-if)#interface po2`	Enters the configuration mode for po2.
`PE5(config-if)#ip address 70.1.1.2/24`	Assigns the IP address 70.1.1.2 with a subnet mask of /24 to this interface.
`PE5(config-if)#load-interval 30`	Sets the load interval to 30 seconds for monitoring the interface.
`PE5(config-if)#interface sa1`	Assigns the IP address 10.1.1.2 with a subnet mask of /24 to this interface.
`PE5(config-if)#ip ospf cost 10`	Sets the OSPF cost for this interface to 10.
`PE5(config-if)#load-interval 30`	Sets the load interval to 30 seconds for monitoring the interface.
`PE5(config-if)#interface ce50`	Enters the configuration mode for ce50.
`PE5(config-if)#ip address 50.1.1.1/24`	Assigns the IP address 50.1.1.1 with a subnet mask of /24 to this interface.
`PE5(config-if)#ip router isis 1`	Specifies that ISIS routing process 1 is enabled on this interface.
`PE5(config-if)#load-interval 30`	Sets the load interval to 30 seconds for monitoring the interface.

`PE5(config-if)#router ospf 1`	Enters ISIS configuration mode with process ID 1.
`PE5(config-if)#network 10.1.1.0/24 area 0.0.0.0`	Specifies that the network 10.1.1.0 with subnet mask 255.255.255.0 belongs to OSPF area 0.0.0.0.
`PE5(config-if)#network 30.1.1.0/24 area 0.0.0.0`	Specifies another network, 30.1.1.0 with subnet mask 255.255.255.0, also belonging to OSPF area 0.0.0.0.
`PE5(config-if)#network 70.1.1.0/24 area 0.0.0.0`	Specifies a third network, 70.1.1.0 with subnet mask 255.255.255.0, in OSPF area 0.0.0.0.
`PE5(config-if)#router isis 1`	Enters ISIS configuration mode with process ID 1.
`PE5(config-if)#is-type level-1-2`	Configures this ISIS router to support both Level 1 and Level 2 routing.
`PE5(config-if)#metric-style wide`	Configures ISIS to use the wide metric style, which allows for greater flexibility in metric values.
`PE5(config-if)# mpls traffic-eng router-id 5.5.5.5`	Sets the MPLS Traffic Engineering router ID to 5.5.5.5.
`PE5(config-if)#mpls traffic-eng level-1`	Enables MPLS Traffic Engineering for Level 1 routing.
`PE5(config-if)#mpls traffic-eng level-2`	Enables MPLS Traffic Engineering for Level 2 routing.
`PE5(config-if)#dynamic-hostname`	Allows the hostname to be dynamically generated.
`PE5(config-if)#bfd all-interfaces`	Enables Bidirectional Forwarding Detection on all interfaces.
`PE5(config-if)#net 49.0000.0005.0001.00`	Sets the NET for this router.
`PE5(config-if)#exit`	Exits from the router mode.

PE 6

`PE6#configure terminal`	Enters the configuration maode.
`PE6(config-if)#interface sa2`	Enters configuration mode for interface sa2.
`PE6(config-if)#ip address 80.1.1.2/24`	Assigns the IP address 80.1.1.2 with a subnet mask of 255.255.255.0 to interface sa2.
`PE6(config-if)#ip router isis 1`	Associates ISIS routing protocol with this interface using process ID 1.
`PE6(config-if)#load-interval 30`	Sets the load-interval to 30 seconds.
`PE6(config-if)#interface ce1`	Enters configuration mode for interface ce1.
`PE6(config-if)#ip address 50.1.1.2/24`	Assigns the IP address 50.1.1.2 with a subnet mask of 255.255.255.0 to interface ce1.
`PE6(config-if)#ip router isis 1`	Associates ISIS routing protocol with this interface using process ID 1.
`PE6(config-if)#load-interval 30`	Sets the load-interval to 30 seconds.
`PE6(config-if)#interface ce2`	Enters configuration mode for interface ce2.
`PE6(config-if)#speed 40g`	Sets the interface speed to 40 gigabits per second.
`PE6(config-if)#ip address 40.1.1.1/24`	Assigns the IP address 40.1.1.1 with a subnet mask of 255.255.255.0 to interface ce2.
`PE6(config-if)#ip ospf cost 10`	Sets the OSPF cost for this interface to 10.
`PE6(config-if)#load-interval 30`	Sets the load-interval to 30 seconds.
`PE6(config-if)#router ospf 1`	Enters ISIS configuration mode with process ID 1.
`PE6(config-if)#network 30.1.1.0/24 area 0.0.0.0`	Specifies another network, 30.1.1.0 with subnet mask 255.255.255.0, also belonging to OSPF area 0.0.0.0.

`PE6(config-if)#network 40.1.1.0/24 area 0.0.0.0`	Specifies a third network, 40.1.1.0/24 with subnet mask 255.255.255.0, in OSPF area 0.0.0.0.
`PE6(config-if)#router isis 1`	Enters ISIS configuration mode with process ID 1.
`PE6(config-if)#is-type level-1-2`	Configures this ISIS router to support both Level 1 and Level 2 routing.
`PE6(config-if)#metric-style wide`	Configures ISIS to use the wide metric style, which allows for greater flexibility in metric values.
`PE6(config-if)# mpls traffic-eng router-id 6.6.6.6`	Sets the MPLS Traffic Engineering router ID to 6.6.6.6.
`PE6(config-if)#mpls traffic-eng level-1`	Enables MPLS Traffic Engineering for Level 1 routing.
`PE6(config-if)#mpls traffic-eng level-2`	Enables MPLS Traffic Engineering for Level 2 routing.
`PE6(config-if)#dynamic-hostname`	Allows the hostname to be dynamically generated.
`PE6(config-if)#bfd all-interfaces`	Enable BFD on all network interfaces.

PE4

`PE4#configure terminal`	Enters the configuration mode.
`PE4(config-if)# interface xe5`	Enters configuration mode for xe5.
`PE4(config-if)# ip address 60.1.1.2/24`	Assigns the IP address 60.1.1.2 with a subnet mask of 255.255.255.0 to the interface.
`PE4(config-if)#ip router isis 1`	Enables ISIS routing protocol on the interface with process ID 1.
`PE4(config-if)#load-interval 30`	Sets the interval for which interface statistics are collected to 30 seconds.
`PE4(config)#nvo vxlan enable`	Enables the VXLAN feature on the device.
`PE4(config)#nvo vxlan irb`	Enables VXLAN IRB functionality.
`PE4(config-vrf)#mac vrf L2VRF2`	Configures a VRF instance named L2VRF2 and associates it with a specific RD
`PE4(config-vrf)# rd 3.3.3.3:11`	Sets the RD for the L2VRF2 VRF to 3.3.3.3:11.
`PE4(config-vrf)#route-target both 10.10.10.10:100`	Associates a route target with the L2VRF2 VRF for VPN route distribution.
`PE4(config-vrf)#ip vrf L3VRF2`	Configures another VRF named L3VRF2.
`PE4(config-vrf)#rd 63000:11`	Sets the RD for the L3VRF2 VRF to 63000:11.
`PE4(config-vrf)# route-target both 101:101`	Associates a route target with the L3VRF2 VRF for VPN route distribution.
`PE4(config-vrf)# l3vni 2000`	Configures the L3VNI for the L3VRF2 VRF.
`PE4(config)#interface irb2001`	Configuring an IRB interface with the number 2001.
`PE4(config-irb-if)# ip vrf forwarding L3VRF2`	Associates the IRB interface with the L3VRF2 VRF.
`PE4(config-irb-if)# ip address 14.14.14.1/24`	Assigns an IP address to the IRB interface.
`PE4(config-irb-if)#mtu 9000`	Sets the MTU for the IRB interface.
`PE4(config-irb-if)#ip router isis 2`	Associates the IRB interface with ISIS routing.
`PE4(config-irb)#interface irb3002`	Configures another IRB interface with the number 3002.
`PE4(config-irb-if)# ip vrf forwarding L3VRF2`	Associates the second IRB interface with the "L3VRF2" VRF.
`PE4(config-irb-if)# ipv6 address 3002::1/64`	Assigns an IPv6 address to the second IRB interface.
`PE4(config-irb-if)#mtu 9000`	Sets the MTU for the second IRB interface.

`PE4(config-irb)#ipv6 router isis 3`	Associates the IRB interfaces with IPv6 and ISIS routing.
`PE4(config)#router isis 2 L3VRF2`	Configures ISIS routing with the VRF L3VRF2.
`PE4(config-router)#is-type level-1-2`	Sets the ISIS level type to level-1-2.
`PE4(config-router)# metric-style wide`	Configures a wide metric style for ISIS.
`PE4(config-router)# dynamic-hostname`	Enables dynamic hostname assignment for the ISIS router.
`PE4(config-router)#bfd all-interfaces`	Enables BFD on all interfaces within ISIS.
`PE4(config-router)# net 49.0000.0000.0441.00`	Configures the network entity title (NET) for ISIS routing with the specified value.
`PE4(config)#router isis 3 L3VRF2`	Configures ISIS routing with the VRF L3VRF2.
`PE4(config-router)#is-type level-1-2`	Sets the ISIS level type to level-1-2.
`PE4(config-router)# metric-style wide`	Configures a wide metric style for ISIS.
`PE4(config-router)# dynamic-hostname`	Enables dynamic hostname assignment for the ISIS router.
`PE4(config-router)#bfd all-interfaces`	Enables BFD on all interfaces within ISIS.
`PE4(config-router)# net 49.0000.0000.0442.00`	Configures the network entity title (NET) for ISIS routing with the specified value.
`PE4(config)#nvo vxlan vtep-ip-global 3.3.3.3`	Configures the global VxLAN VTEP IP address to 3.3.3.3.
`PE4(config)#nvo vxlan id 201 ingress-replication`	Configures a VxLAN with VNI 201 and specifies ingress-replication for multicast traffic handling.
`PE4(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Specifies the EVPN-BGP host-reachability-protocol for the VxLAN with the VRF L2VRF2
`PE4(config-nvo)# evpn irb2001`	Enables EVPN IRB (Integrated Routing and Bridging) for VxLAN interface IRB2001.
`PE4(config-nvo)# vni-name VNI-201`	Assigns a name VNI-201 to the VxLAN VNI 201.
`PE4(config)#nvo vxlan id 3002 ingress-replication`	Configures another VxLAN with VNI 3002 and specifies ingress-replicationr for multicast traffic handling.
`PE4(config-nvo)# vxlan host-reachability-protocol evpn-bgp L2VRF2`	Specifies the EVPN-BGP host-reachability-protocol for the VxLAN with the VRF L2VRF2.
`PE4(config-nvo)# evpn irb3002`	Enables EVPN IRB for VxLAN interface IRB3002
`PE4(config-if)#interface xe5`	Enters the configuration mode for the interface 5.
`PE4(config-if)#switchport`	Configures the interface as a L2 switchport.
`PE4(config-if)#load-interval 30`	Sets the interval for which interface statistics are collected to 30 seconds.
`PE4(config)#nvo vxlan access-if port-vlan xe5 100`	Configures a VxLAN network virtualization overlay (NVO) on the interface xe2 with VLAN ID 100
`PE4(config-nvo-acc-if)# map vnid 101`	Maps VLAN 100 to the VxLAN Network Identifier (VNID) 101.
`PE4(config-nvo-acc-if)#nvo vxlan access-if port-vlan xe5 2001`	Configures another VxLAN NVO on the same interface xe2, but this time with VLAN ID 2001
`PE4(config-nvo-acc-if)# map vnid 2001`	Maps VLAN 2001 to a different VxLAN VNID.
`PE4(config-if)#router isis 1`	Starts the ISIS routing process with process ID 1.
`PE4(config-if)#is-type level-1-2`	Specifies that the router participates in both Level 1 and Level 2 routing.
`PE4(config-if)#metric-style wide`	Configures the metric style to be wide, enabling more flexibility in metric calculations.
`PE4(config-if)#mpls traffic-eng router-id 2.2.2.2`	Sets the MPLS Traffic Engineering router ID to 2.2.2.2.

PE4(config-if)#mpls traffic-eng level-1	Enables MPLS Traffic Engineering for Level 1 ISIS.
PE4(config-if)#mpls traffic-eng level-2	Enables MPLS Traffic Engineering for Level 2 ISIS.
PE4(config-if)#dynamic-hostname	Enables the dynamic hostname feature for ISIS.
PE4(config-if)#bfd all-interfaces	Configures Bidirectional Forwarding Detection on all interfaces.
PE4(config-if)#net 49.0000.0003.0001.00	Specifies the network entity title (NET) for ISIS.

BGP Configuration

PE4(config)#router bgp 100	Starts the BGP routing process with an autonomous system number (AS) of 100.
PE4(config-router)#bgp router-id 3.3.3.3	Sets the BGP router ID to 3.3.3.3
PE4(config-router)#neighbor 2.2.2.2 remote-as 100	Configures a BGP neighbor with the IP address 2.2.2.2 and specifies the remote AS number as 100.
PE4(config-router)#neighbor 2.2.2.2 update-source lo	Specifies that loopback interface (lo) is the source for BGP updates to the neighbor.
PE4(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Sets the advertisement interval to 0, which means updates will be sent immediately.
PE4(config-router)#address-family ipv4 unicast	Enters the configuration mode for the IPv4 unicast address family within the router configuration.
PE4(config-router-af)#network 3.3.3.3/32	Specifies that network 3.3.3.3 with a /32 subnet mask is part of the IPv4 unicast address family.
PE4(config-router-af)#neighbor 2.2.2.2 activate	Activates the neighbor with the IP address 2.2.2.2 for the IPv4 unicast address family.
PE4(config-router-af)#exit-address-family	Exits the configuration mode for the IPv4 unicast address family.
PE4(config-router)#address-family l2vpn evpn	Enters the configuration mode for the L2VPN EVPN address family within the router configuration.
PE4(config-router-af)#neighbor 2.2.2.2 activate	Activates the neighbor with the IP address 2.2.2.2 for the L2VPN EVPN address family.
PE4(config-router-af)#exit-address-family	Exits the configuration mode for the L2VPN EVPN address family.
PE4(config-router)#address-family ipv4 vrf L3VRF2	Enters the configuration mode for the IPv4 address family within the VRF named L3VRF2.
PE4(config-router-af)#redistribute connected	Configures the redistribution of directly connected routes into the IPv4 address family for the specified VRF.
PE4(config-router-af)#exit-address-family	Exits the configuration mode for the IPv4 address family within the VRF L3VRF2.
PE4(config-router-af)#address-family ipv6 vrf L3VRF2	Enters the configuration mode for the IPv6 address family within the VRF named L3VRF2.
PE4(config-router-af)#redistribute connected	Configures the redistribution of directly connected routes into the IPv6 address family for the specified VRF.
PE4(config-router-af)#exit-address-family	Exits the configuration mode for the IPv6 address family within the VRF L3VRF2.

Implementation Examples

Scenario: Configure OSPF and ISIS protocols on an IRB interface with an assigned IP address.

New CLI Commands

No CLI commands are introduced.

Validation

OSPF Validation

```
PE1#show ip ospf neighbor
Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
       Instance ID
50.1.1.1          1   Full/DR         00:00:38    10.1.1.2        sa1
           0

Total number of full neighbors: 1
OSPF process 2 VRF(L3VRF1):
Neighbor ID     Pri   State           Dead Time   Address         Interface
       Instance ID
192.0.0.1         0   Full/DROther    00:00:34    11.11.11.2      irb1001
           0
PE1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
1.1.1.1          4.4.4.4          Installed     00:15:59      00:15:59

Total number of entries are 2
PE1#show nvo vxlan irb-status
IRB is ACTIVE in Hardware
PE1#show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
________________________________________________________________________________
101      11.11.11.1       9819.2ccd.9301 Static Local   ----
101      11.11.11.2       0010.9400.0001 Dynamic Local  ----
Total number of entries are 2
PE1#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default
```

```
IP Route Table for VRF "default"
C            1.1.1.1/32 is directly connected, lo, 00:53:03
O            4.4.4.4/32 [110/31] via 10.1.1.2, sa1, 00:16:29
O            7.7.7.7/32 [110/12] via 10.1.1.2, sa1, 00:44:26
C            10.1.1.0/24 is directly connected, sa1, 00:50:10
O            30.1.1.0/24 [110/20] via 10.1.1.2, sa1, 00:44:22
O            40.1.1.0/24 [110/30] via 10.1.1.2, sa1, 00:17:14
O            70.1.1.0/24 [110/11] via 10.1.1.2, sa1, 00:45:18
C            127.0.0.0/8 is directly connected, lo, 00:53:03
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 00:53:03
C            127.0.0.0/8 is directly connected, lo.management, 00:53:03
IP Route Table for VRF "L2VRF1"
IP Route Table for VRF "L3VRF1"
B            4.4.4.4/32 [0/0] is directly connected, tunvxlan2, 00:16:25
B            7.7.7.7/32 [0/0] is directly connected, tunvxlan2, 00:44:21
C            11.11.11.0/24 is directly connected, irb1001, 00:53:03
B            12.12.12.0/24 [200/0] via 4.4.4.4 (recursive is directly connected,
 tunvxlan2), 00:16:26
B            16.16.16.0/24 [200/0] via 7.7.7.7 (recursive is directly connected,
 tunvxlan2), 00:44:21
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:53:03

Gateway of last resort is not set


PE1#show bgp l2vpn evpn
BGP table version is 5, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i
- internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop          Metric    LocPrf    Weight     Path Peer       Encap

RD[7100:11]
*>i  [5]:[0]:[0]:[24]:[16.16.16.0]:[0.0.0.0]:[1000]
                      7.7.7.7             0        100       0    i  7.7.7.7       VXLAN
*>i  [5]:[0]:[0]:[64]:[7002::]:[::]:[1000]
                      7.7.7.7             0        100       0    i  7.7.7.7       VXLAN

RD[56000:11]
```

```
*>i  [5]:[0]:[0]:[24]:[12.12.12.0]:[0.0.0.0]:[1000]
                     4.4.4.4             0        100       0    ?  4.4.4.4       VXLAN
*>i  [5]:[0]:[0]:[64]:[2002::]:[::]:[1000]
                     4.4.4.4             0        100       0    ?  4.4.4.4       VXLAN

RD[1.1.1.1:11] VRF[L2VRF1]:
*>   [2]:[0]:[101]:[48,0010:9400:0001]:[0]:[101]
                    1.1.1.1              0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[101]:[48,0010:9400:0001]:[32,11.11.11.2]:[101]
                    1.1.1.1              0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[101]:[48,9819:2ccd:9301]:[32,11.11.11.1]:[101]
                    1.1.1.1              0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[102]:[48,0010:9400:0002]:[0]:[102]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
* i  [2]:[0]:[102]:[48,0010:9400:0002]:[32,12.12.12.2]:[102]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
* i  [2]:[0]:[102]:[48,5c07:5813:425e]:[32,12.12.12.1]:[102]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
*>   [2]:[0]:[2001]:[48,0010:9400:0009]:[0]:[2001]
                    1.1.1.1              0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[2001]:[48,0010:9400:0009]:[128,2001::2][2001]
                    1.1.1.1              0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[2001]:[48,9819:2ccd:9301]:[128,2001::1][2001]
                    1.1.1.1              0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[2002]:[48,0010:9400:000a]:[0]:[2002]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
* i  [2]:[0]:[2002]:[48,0010:9400:000a]:[128,2002::2][2002]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
* i  [2]:[0]:[2002]:[48,5c07:5813:425e]:[128,2002::1][2002]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
*>   [3]:[101]:[32,1.1.1.1]
                    1.1.1.1              0       100      32768  i  ---------      VXLAN
* i  [3]:[102]:[32,4.4.4.4]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
*>   [3]:[2001]:[32,1.1.1.1]
                    1.1.1.1              0       100      32768  i  ---------      VXLAN
* i  [3]:[2002]:[32,4.4.4.4]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN

RD[4.4.4.4:11]
*>i  [2]:[0]:[102]:[48,0010:9400:0002]:[0]:[102]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
*>i  [2]:[0]:[102]:[48,0010:9400:0002]:[32,12.12.12.2]:[102]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
*>i  [2]:[0]:[102]:[48,5c07:5813:425e]:[32,12.12.12.1]:[102]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
*>i  [2]:[0]:[2002]:[48,0010:9400:000a]:[0]:[2002]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
*>i  [2]:[0]:[2002]:[48,0010:9400:000a]:[128,2002::2][2002]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
```

```
*>i  [2]:[0]:[2002]:[48,5c07:5813:425e]:[128,2002::1][2002]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
*>i  [3]:[102]:[32,4.4.4.4]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN
*>i  [3]:[2002]:[32,4.4.4.4]
                     4.4.4.4             0        100       0    i  4.4.4.4       VXLAN

Total number of prefixes 28

PE3#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=================================================================================
4.4.4.4          1.1.1.1          Installed     00:18:19      00:18:19

Total number of entries are 1
PE3#show ip ospf neighbor

Total number of full neighbors: 1
OSPF process 1 VRF(default):
Neighbor ID     Pri   State            Dead Time   Address         Interface
       Instance ID
40.1.1.2          1   Full/DR          00:00:36    40.1.1.1        ce30
           0

Total number of full neighbors: 1
OSPF process 2 VRF(L3VRF1):
Neighbor ID     Pri   State            Dead Time   Address         Interface
       Instance ID
192.0.0.2         0   Full/DROther     00:00:36    12.12.12.2      irb1001
           0
PE3#show ip route vrf all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

IP Route Table for VRF "default"
O            1.1.1.1/32 [110/31] via 40.1.1.1, ce30, 00:18:35
C            4.4.4.4/32 is directly connected, lo, 00:19:22
O            7.7.7.7/32 [110/22] via 40.1.1.1, ce30, 00:18:35
O            10.1.1.0/24 [110/30] via 40.1.1.1, ce30, 00:18:35
O            30.1.1.0/24 [110/20] via 40.1.1.1, ce30, 00:18:35
C            40.1.1.0/24 is directly connected, ce30, 00:19:21
O            70.1.1.0/24 [110/21] via 40.1.1.1, ce30, 00:18:35
C            127.0.0.0/8 is directly connected, lo, 00:20:05
```

```
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 00:19:19
C            127.0.0.0/8 is directly connected, lo.management, 00:20:05
IP Route Table for VRF "L3VRF1"
B            1.1.1.1/32 [0/0] is directly connected, tunvxlan2, 00:18:31
B            11.11.11.0/24 [200/0] via 1.1.1.1 (recursive is directly connected,
 tunvxlan2), 00:18:32
C            12.12.12.0/24 is directly connected, irb1001, 00:19:28
C            127.0.0.0/8 is directly connected, lo.L3VRF1, 00:19:29
IP Route Table for VRF "L2VRF1"

Gateway of last resort is not set
PE3# show bgp l2vpn evpn
BGP table version is 4, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i
- internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop       Metric    LocPrf     Weight  Path Peer      Encap

RD[51000:11]
*>i  [5]:[0]:[0]:[24]:[11.11.11.0]:[0.0.0.0]:[1000]
                     1.1.1.1             0        100       0    ?  1.1.1.1       VXLAN
*>i  [5]:[0]:[0]:[64]:[2001::]:[::]:[1000]
                     1.1.1.1             0        100       0    ?  1.1.1.1       VXLAN

RD[1.1.1.1:11]
*>i  [2]:[0]:[101]:[48,0010:9400:0001]:[0]:[101]
                     1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[101]:[48,0010:9400:0001]:[32,11.11.11.2]:[101]
                     1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[101]:[48,9819:2ccd:9301]:[32,11.11.11.1]:[101]
                     1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[2001]:[48,0010:9400:0009]:[0]:[2001]
                     1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[2001]:[48,0010:9400:0009]:[128,2001::2][2001]
                     1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [2]:[0]:[2001]:[48,9819:2ccd:9301]:[128,2001::1][2001]
                     1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [3]:[101]:[32,1.1.1.1]
                     1.1.1.1             0        100       0    i  1.1.1.1       VXLAN
*>i  [3]:[2001]:[32,1.1.1.1]
```

```
                      1.1.1.1              0        100       0    i  1.1.1.1       VXLAN

RD[4.4.4.4:11] VRF[L2VRF1]:
* i  [2]:[0]:[101]:[48,0010:9400:0001]:[0]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
* i  [2]:[0]:[101]:[48,0010:9400:0001]:[32,11.11.11.2]:[101]
                      1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
* i  [2]:[0]:[101]:[48,9819:2ccd:9301]:[32,11.11.11.1]:[101]
                       1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
*>   [2]:[0]:[102]:[48,0010:9400:0002]:[0]:[102]
                     4.4.4.4              0        100      32768  i  ---------      VXLAN
*>   [2]:[0]:[102]:[48,0010:9400:0002]:[32,12.12.12.2]:[102]
                     4.4.4.4              0        100      32768  i  ---------      VXLAN
*>   [2]:[0]:[102]:[48,5c07:5813:425e]:[32,12.12.12.1]:[102]
                      4.4.4.4              0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[2001]:[48,0010:9400:0009]:[0]:[2001]
                      1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
* i  [2]:[0]:[2001]:[48,0010:9400:0009]:[128,2001::2][2001]
                      1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
* i  [2]:[0]:[2001]:[48,9819:2ccd:9301]:[128,2001::1][2001]
                       1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
*>   [2]:[0]:[2002]:[48,0010:9400:000a]:[0]:[2002]
                     4.4.4.4              0        100      32768  i  ---------      VXLAN
*>   [2]:[0]:[2002]:[48,0010:9400:000a]:[128,2002::2][2002]
                     4.4.4.4              0        100      32768  i  ---------      VXLAN
*>   [2]:[0]:[2002]:[48,5c07:5813:425e]:[128,2002::1][2002]
                     4.4.4.4              0        100      32768  i  ---------      VXLAN
* i  [3]:[101]:[32,1.1.1.1]
                       1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
*>   [3]:[102]:[32,4.4.4.4]
                     4.4.4.4              0        100      32768  i  ---------      VXLAN
* i  [3]:[2001]:[32,1.1.1.1]
                       1.1.1.1              0        100       0    i  1.1.1.1       VXLAN
*>   [3]:[2002]:[32,4.4.4.4]
                     4.4.4.4              0        100      32768  i  ---------      VXLAN

Total number of prefixes 26
```

ISIS Validation

```
PE2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
2.2.2.2          3.3.3.3          Installed     00:00:10      00:00:10

Total number of entries are 1
PE2#show clns neighbors

Total number of L1 adjacencies: 1
```

```
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE5            po1         b86a.9725.a7f2  Up     28        L1   IS-IS
                                           Up     28        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 2:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb2001     0010.9400.0003  Up     28        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 3:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb3001     0010.9400.000c  Up     28        L2   IS-IS
PE2#show ip route vrf all
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            2.2.2.2/32 is directly connected, lo, 02:13:57
i L2         3.3.3.3/32 [115/30] via 20.1.1.2, po1, 00:00:32
i L1         7.7.7.7/32 [115/40] via 20.1.1.2, po1, 01:05:49
C            20.1.1.0/24 is directly connected, po1, 02:13:21
i L1         50.1.1.0/24 [115/20] via 20.1.1.2, po1, 01:06:05
i L1         60.1.1.0/24 [115/30] via 20.1.1.2, po1, 00:00:47
i L1         80.1.1.0/24 [115/30] via 20.1.1.2, po1, 01:05:49
C            127.0.0.0/8 is directly connected, lo, 02:13:57
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 02:13:57
C            127.0.0.0/8 is directly connected, lo.management, 02:13:57
IP Route Table for VRF "L3VRF2"
B            3.3.3.3/32 [0/0] is directly connected, tunvxlan2, 00:00:28
C            13.13.13.0/24 is directly connected, irb2001, 02:13:57
B            14.14.14.0/24 [200/0] via 3.3.3.3 (recursive is directly connected,
 tunvxlan2), 00:00:28
C            127.0.0.0/8 is directly connected, lo.L3VRF2, 02:13:57
IP Route Table for VRF "L2VRF2"
```

```
Gateway of last resort is not set
PE2# show bgp l2vpn evpn
BGP table version is 2, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i
- internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop          Metric    LocPrf    Weight   Path Peer         Encap

RD[63000:11]
*>i  [5]:[0]:[0]:[24]:[14.14.14.0]:[0.0.0.0]:[2000]
                       3.3.3.3             0        100       0    ?  3.3.3.3       VXLAN
*>i  [5]:[0]:[0]:[64]:[3002::]:[::]:[2000]
                       3.3.3.3             0        100       0    ?  3.3.3.3       VXLAN

RD[2.2.2.2:11] VRF[L2VRF2]:
*>   [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0        100      32768  i  ---------      VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0        100      32768  i  ---------      VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                       3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                       3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2             0        100      32768  i  ---------      VXLAN
* i  [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                       3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0        100      32768  i  ---------      VXLAN
*>   [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0        100      32768  i  ---------      VXLAN
*>   [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0        100      32768  i  ---------      VXLAN
* i  [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                       3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                       3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                       3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0        100      32768  i  ---------      VXLAN
```

```
* i  [3]:[201]:[32,3.3.3.3]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>   [3]:[3001]:[32,2.2.2.2]
                    2.2.2.2             0        100      32768  i  ---------      VXLAN
* i  [3]:[3002]:[32,3.3.3.3]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN

RD[3.3.3.3:11]
*>i  [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [3]:[201]:[32,3.3.3.3]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [3]:[3002]:[32,3.3.3.3]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN

Total number of prefixes 26
PE2# show nvo vxlan arp-
arp-cache  arp-nd
PE2# show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
============================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
201      13.13.13.1       e8c5.7a76.581d Static Local   ----
201      13.13.13.2       0010.9400.0003 Dynamic Local  ----
201      14.14.14.1       e8c5.7aa8.7cb3 Static Remote  ----
201      14.14.14.2       0010.9400.0005 Dynamic Remote ----
Total number of entries are 4
PE2#show nvo vxlan irb-status
IRB is ACTIVE in Hardware
PE2#

PE4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=============================================================================
3.3.3.3          7.7.7.7          Installed     00:01:28      00:01:28
3.3.3.3          2.2.2.2          Installed     00:01:28      00:01:28

Total number of entries are 2
```

```
PE4#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE6            xe5         00e0.4b71.f12c  Up     25        L1   IS-IS
                                           Up     25        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 2:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb2001     0010.9400.0005  Up     28        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 3:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb3002     0010.9400.000b  Up     28        L2   IS-IS
PE4#show ip route vrf all
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 L2         2.2.2.2/32 [115/30] via 60.1.1.1, xe5, 00:01:46
C            3.3.3.3/32 is directly connected, lo, 02:09:52
i L1         7.7.7.7/32 [115/30] via 60.1.1.1, xe5, 00:01:46
i L1         20.1.1.0/24 [115/30] via 60.1.1.1, xe5, 00:01:46
i L1         50.1.1.0/24 [115/20] via 60.1.1.1, xe5, 00:01:46
C            60.1.1.0/24 is directly connected, xe5, 00:02:02
i L1         80.1.1.0/24 [115/20] via 60.1.1.1, xe5, 00:01:46
C            127.0.0.0/8 is directly connected, lo, 02:09:52
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 02:09:52
C            127.0.0.0/8 is directly connected, lo.management, 02:09:52
IP Route Table for VRF "L3VRF2"
B            2.2.2.2/32 [0/0] is directly connected, tunvxlan2, 00:01:42
B            7.7.7.7/32 [0/0] is directly connected, tunvxlan2, 00:01:42
B            13.13.13.0/24 [200/0] via 2.2.2.2 (recursive is directly connected,
tunvxlan2), 00:01:42
C            14.14.14.0/24 is directly connected, irb2001, 02:09:52
```

```
B           17.17.17.0/24 [200/0] via 7.7.7.7 (recursive is directly connected,
tunvxlan2), 00:01:42
C           127.0.0.0/8 is directly connected, lo.L3VRF2, 02:09:52
IP Route Table for VRF "L2VRF2"

Gateway of last resort is not set
PE4# show bgp l2vpn evpn
BGP table version is 3, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

    Network          Next Hop            Metric    LocPrf       Weight     Path  Peer  Encap

RD[7400:11]
*>i  [5]:[0]:[0]:[24]:[17.17.17.0]:[0.0.0.0]:[2000]
                     7.7.7.7             0         100        0     i  7.7.7.7         VXLAN
*>i  [5]:[0]:[0]:[64]:[8002::]:[::]:[2000]
                     7.7.7.7             0         100        0     i  7.7.7.7         VXLAN

RD[61000:11]
*>i  [5]:[0]:[0]:[24]:[13.13.13.0]:[0.0.0.0]:[2000]
                     2.2.2.2             0         100        0     ?  2.2.2.2         VXLAN
*>i  [5]:[0]:[0]:[64]:[3001::]:[::]:[2000]
                     2.2.2.2             0         100        0     ?  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0         100        0     i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0         100        0     i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2             0         100        0     i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0         100        0     i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0         100        0     i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0         100        0     i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0         100        0     i  2.2.2.2         VXLAN
*>i  [3]:[3001]:[32,2.2.2.2]
                     2.2.2.2             0         100        0     i  2.2.2.2         VXLAN
```

```
RD[3.3.3.3:11] VRF[L2VRF2]:
* i  [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                      2.2.2.2             0       100       0    i  2.2.2.2        VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                      2.2.2.2             0       100       0    i  2.2.2.2        VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                     3.3.3.3             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                      3.3.3.3              0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                      2.2.2.2             0       100       0    i  2.2.2.2        VXLAN
*>   [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                      3.3.3.3              0        100       32768  i  ----------
VXLAN
* i  [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                      2.2.2.2             0       100       0    i  2.2.2.2        VXLAN
* i  [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                      2.2.2.2             0       100       0    i  2.2.2.2        VXLAN
* i  [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                      2.2.2.2             0       100       0    i  2.2.2.2        VXLAN
*>   [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                     3.3.3.3             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                     3.3.3.3             0       100      32768  i  ----------     VXLAN
*>   [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                     3.3.3.3             0       100      32768  i  ----------     VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2        VXLAN
*>   [3]:[201]:[32,3.3.3.3]
                     3.3.3.3             0       100      32768  i  ---------      VXLAN
* i  [3]:[3001]:[32,2.2.2.2]
                      2.2.2.2              0        100       0    i  2.2.2.2        VXLAN
*>   [3]:[3002]:[32,3.3.3.3]
                     3.3.3.3             0       100      32768  i  ----------     VXLAN

Total number of prefixes 28
```

ISIS Validation

```
PE2#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
2.2.2.2          3.3.3.3          Installed     00:00:10      00:00:10

Total number of entries are 1
PE2#show clns neighbors
Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
```

```
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE5            po1         b86a.9725.a7f2  Up     28        L1   IS-IS
                                           Up     28        L2   IS-IS
Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 2:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb2001     0010.9400.0003  Up     28        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 3:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb3001     0010.9400.000c  Up     28        L2   IS-IS
PE2#
PE2#
PE2#show ip route vrf all
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            2.2.2.2/32 is directly connected, lo, 02:13:57
i L2         3.3.3.3/32 [115/30] via 20.1.1.2, po1, 00:00:32
i L1         7.7.7.7/32 [115/40] via 20.1.1.2, po1, 01:05:49
C            20.1.1.0/24 is directly connected, po1, 02:13:21
i L1         50.1.1.0/24 [115/20] via 20.1.1.2, po1, 01:06:05
i L1         60.1.1.0/24 [115/30] via 20.1.1.2, po1, 00:00:47
i L1         80.1.1.0/24 [115/30] via 20.1.1.2, po1, 01:05:49
C            127.0.0.0/8 is directly connected, lo, 02:13:57
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 02:13:57
C            127.0.0.0/8 is directly connected, lo.management, 02:13:57
IP Route Table for VRF "L3VRF2"
B            3.3.3.3/32 [0/0] is directly connected, tunvxlan2, 00:00:28
C            13.13.13.0/24 is directly connected, irb2001, 02:13:57
B            14.14.14.0/24 [200/0] via 3.3.3.3 (recursive is directly connected,
 tunvxlan2), 00:00:28
C            127.0.0.0/8 is directly connected, lo.L3VRF2, 02:13:57
IP Route Table for VRF "L2VRF2"
```

```
Gateway of last resort is not set
PE2# show bgp l2vpn evpn
BGP table version is 2, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i
- internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route
Network          Next Hop            Metric    LocPrf       Weight     Path
 Peer          Encap

RD[63000:11]
*>i  [5]:[0]:[0]:[24]:[14.14.14.0]:[0.0.0.0]:[2000]
                      3.3.3.3             0        100       0    ?  3.3.3.3       VXLAN
*>i  [5]:[0]:[0]:[64]:[3002::]:[::]:[2000]
                      3.3.3.3             0        100       0    ?  3.3.3.3       VXLAN

RD[2.2.2.2:11] VRF[L2VRF2]:
*>   [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2              0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
*>   [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
* i  [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
* i  [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                      3.3.3.3             0        100       0    i  3.3.3.3       VXLAN
*>   [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0       100      32768  i  ---------      VXLAN
```

```
* i  [3]:[201]:[32,3.3.3.3]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>   [3]:[3001]:[32,2.2.2.2]
                    2.2.2.2             0        100       32768  i  ---------     VXLAN
* i  [3]:[3002]:[32,3.3.3.3]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN

RD[3.3.3.3:11]
*>i  [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [3]:[201]:[32,3.3.3.3]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN
*>i  [3]:[3002]:[32,3.3.3.3]
                     3.3.3.3              0        100       0    i  3.3.3.3       VXLAN

Total number of prefixes 26
PE2#        show nvo vxlan arp-
arp-cache  arp-nd
PE2#        show nvo vxlan arp-cache
VXLAN ARP-CACHE Information
===========================
VNID     Ip-Addr          Mac-Addr       Type        Age-Out    Retries-Left
_____________________________________________________________________________
201      13.13.13.1       e8c5.7a76.581d Static Local   ----
201      13.13.13.2       0010.9400.0003 Dynamic Local  ----
201      14.14.14.1       e8c5.7aa8.7cb3 Static Remote  ----
201      14.14.14.2       0010.9400.0005 Dynamic Remote ----
Total number of entries are 4
PE2#show nvo vxlan irb-status
IRB is ACTIVE in Hardware
PE2#

PE4#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
=============================================================================
3.3.3.3          7.7.7.7          Installed     00:01:28      00:01:28
3.3.3.3          2.2.2.2          Installed     00:01:28      00:01:28

Total number of entries are 2
```

```
PE4#show clns neighbors

Total number of L1 adjacencies: 1
Total number of L2 adjacencies: 1
Total number of adjacencies: 2
Tag 1:  VRF : default
System Id      Interface   SNPA            State  Holdtime  Type Protocol
PE6            xe5         00e0.4b71.f12c  Up     25        L1   IS-IS
                                           Up     25        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 2:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb2001     0010.9400.0005  Up     28        L2   IS-IS

Total number of L1 adjacencies: 0
Total number of L2 adjacencies: 1
Total number of adjacencies: 1
Tag 3:  VRF : L3VRF2
System Id      Interface   SNPA            State  Holdtime  Type Protocol
Spirent-1      irb3002     0010.9400.000b  Up     28        L2   IS-IS
PE4#show ip route vrf all
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 L2         2.2.2.2/32 [115/30] via 60.1.1.1, xe5, 00:01:46
C            3.3.3.3/32 is directly connected, lo, 02:09:52
i L1         7.7.7.7/32 [115/30] via 60.1.1.1, xe5, 00:01:46
i L1         20.1.1.0/24 [115/30] via 60.1.1.1, xe5, 00:01:46
i L1         50.1.1.0/24 [115/20] via 60.1.1.1, xe5, 00:01:46
C            60.1.1.0/24 is directly connected, xe5, 00:02:02
i L1         80.1.1.0/24 [115/20] via 60.1.1.1, xe5, 00:01:46
C            127.0.0.0/8 is directly connected, lo, 02:09:52
IP Route Table for VRF "management"
C            10.12.98.0/24 is directly connected, eth0, 02:09:52
C            127.0.0.0/8 is directly connected, lo.management, 02:09:52
IP Route Table for VRF "L3VRF2"
B            2.2.2.2/32 [0/0] is directly connected, tunvxlan2, 00:01:42
B            7.7.7.7/32 [0/0] is directly connected, tunvxlan2, 00:01:42
B            13.13.13.0/24 [200/0] via 2.2.2.2 (recursive is directly connected,
tunvxlan2), 00:01:42
C            14.14.14.0/24 is directly connected, irb2001, 02:09:52
```

```
B           17.17.17.0/24 [200/0] via 7.7.7.7 (recursive is directly connected,
tunvxlan2), 00:01:42
C           127.0.0.0/8 is directly connected, lo.L3VRF2, 02:09:52
IP Route Table for VRF "L2VRF2"

Gateway of last resort is not set
PE4# show bgp l2vpn evpn
BGP table version is 3, local router ID is 3.3.3.3
Status codes: s suppressed, d damped, h history, a add-path, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

   Network         Next Hop           Metric    LocPrf      Weight     Path  Peer
Encap

RD[7400:11]
*>i  [5]:[0]:[0]:[24]:[17.17.17.0]:[0.0.0.0]:[2000]
                     7.7.7.7             0        100       0    i  7.7.7.7         VXLAN
*>i  [5]:[0]:[0]:[64]:[8002::]:[::]:[2000]
                     7.7.7.7             0        100       0    i  7.7.7.7         VXLAN

RD[61000:11]
*>i  [5]:[0]:[0]:[24]:[13.13.13.0]:[0.0.0.0]:[2000]
                     2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN
*>i  [5]:[0]:[0]:[64]:[3001::]:[::]:[2000]
                     2.2.2.2             0        100       0    ?  2.2.2.2         VXLAN

RD[2.2.2.2:11]
*>i  [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>i  [3]:[3001]:[32,2.2.2.2]
```

```
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN

RD[3.3.3.3:11] VRF[L2VRF2]:
* i  [2]:[0]:[201]:[48,0010:9400:0003]:[0]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[201]:[48,0010:9400:0003]:[32,13.13.13.2]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0005]:[0]:[201]
                     3.3.3.3             0        100       32768  i  ----------     VXLAN
*>   [2]:[0]:[201]:[48,0010:9400:0005]:[32,14.14.14.2]:[201]
                     3.3.3.3             0        100       32768  i  ----------     VXLAN
* i  [2]:[0]:[201]:[48,e8c5:7a76:581d]:[32,13.13.13.1]:[201]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[201]:[48,e8c5:7aa8:7cb3]:[32,14.14.14.1]:[201]
                     3.3.3.3             0        100       32768  i  ----------     VXLAN
* i  [2]:[0]:[3001]:[48,0010:9400:000c]:[0]:[3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[3001]:[48,0010:9400:000c]:[128,3001::2][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
* i  [2]:[0]:[3001]:[48,e8c5:7a76:581d]:[128,3001::1][3001]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [2]:[0]:[3002]:[48,0010:9400:000b]:[0]:[3002]
                     3.3.3.3             0        100       32768  i  ----------     VXLAN
*>   [2]:[0]:[3002]:[48,0010:9400:000b]:[128,3002::2][3002]
                     3.3.3.3             0        100       32768  i  ----------     VXLAN
*>   [2]:[0]:[3002]:[48,e8c5:7aa8:7cb3]:[128,3002::1][3002]
                     3.3.3.3             0        100       32768  i  ----------     VXLAN
* i  [3]:[201]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [3]:[201]:[32,3.3.3.3]
                     3.3.3.3             0        100       32768  i  ----------     VXLAN
* i  [3]:[3001]:[32,2.2.2.2]
                     2.2.2.2             0        100       0    i  2.2.2.2         VXLAN
*>   [3]:[3002]:[32,3.3.3.3]
                     3.3.3.3             0        100       32768  i  ----------     VXLAN

Total number of prefixes 28
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
ECMP	Equal-Cost Multipath
EVPN	Ethernet Virtual Private Netwrok
MPLS	Multiprotocol Label Switching

VxLAN	Virtual Extensible LAN
SR	Segment Routing
IRB	Integrated Routing
OSPF	Open Shortest Path First
ISIS	Intermediate System to Intermediate System

Glossary

The following provides definitions for key terms used throughout this document.

Single Home VxLAN	This refers to a Virtual Extensible LAN (VxLAN) deployment where a single data center or network site is connected to a single external network (usually the internet) for connectivity.
IRB	A networking feature that enables the integration of Layer 3 IP routing and Layer 2 MAC address bridging within the same interface, simplifying network management and resource utilization.
OSPF	A dynamic and efficient link-state routing protocol used to determine the best path for data packets in an IP network. It is characterized by rapid convergence and adaptability, making it suitable for large and dynamic networks.
ISIS	A routing protocol designed for scalability and stability in computer networks, commonly used in large Service Provider networks. It provides a robust framework for routing information exchange.
Layer 3 Routing	Network routing operations at the Network Layer (Layer 3) of the OSI model, focusing on routing IP packets between different subnets or networks.
Layer 2 Bridging	Network bridging operations at the Data Link Layer (Layer 2) of the OSI model, handling the forwarding of data frames based on MAC addresses within the same network segment.
EVPN	Ethernet VPN, a technology that provides advanced and efficient methods for Layer 2 and Layer 3 services in Ethernet networks, often used in data centers and service provider environments.

Fall Back Option for RADIUS Authentication

Overview

Currently, the Remote Authentication Dial-In User Service (RADIUS) server authentication fallback to the local authentication server only when the RADIUS server is not reachable.

This behavior is modified in the current release to forward the authentication request to the local authentication server when the RADIUS authentication is failed or not reachable.

Feature Characteristics

The RADIUS authentication mechanism is enhanced to fallback to local authentication server when the user

- is not present on RADIUS server or
- authentication fails from RADIUS server

To implement the above requirements, the existing CLI `aaa authentication login default fallback error local non-existent-user vrf management` is used to enable fallback to local authentication server. This is disabled by default.

Note: For invalid secret key there is no fallback local authentication.
Console authentication is not supported for RADIUS.

Benefits

By default, the fallback to local authentication is applied when the RADIUS server is unreachable. For other scenarios, enable the fallback using the CLI.

Configuration

Below is the existing CLI used to enable the fallback local authentication server.

```
aaa authentication login default fallback error local non-existent-user vrf
  management
```

Refer to *Authentication, Authorization and Accounting* section in the OcNOS System Management Configuration Guide, Release 6.4.1.

Validation

Configure `aaa authentication` console and verify console authentication:

```
OcNOS#con t
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#radius-server login host 1.1.1.2 seq-num 1 key 0 kumar
OcNOS(config)#commit
OcNOS(config)#aaa authentication login console group radius
OcNOS(config)#commit
OcNOS(config)#exit
```

```
OcNOS#exit

OcNOS#show users
Current user          : (*).  Lock acquired by user : (#).
CLI user              : [C].  Netconf users         : [N].
Location : Applicable to CLI users.
Session  : Applicable to NETCONF users.

Line       User          Idle        Location/Session PID    TYPE   Role
(*) 0 con 0 [C]ocnos      0d00h00m    ttyS0           5531   Remote network-admin
```

Enabled RADIUS local fallback and verify the authentication:

```
OcNOS(config)#aaa authentication login console group radius local
OcNOS(config)#commit
OcNOS(config)#exit
OcNOS#exit
OcNOS>exit

OcNOS>enable
OcNOS#show users
Current user          : (*).  Lock acquired by user : (#).
CLI user              : [C].  Netconf users         : [N].
Location : Applicable to CLI users.
Session  : Applicable to NETCONF users.

Line         User          Idle         Location/Session PID   TYPE   Role
(*) 0 con 0  [C]test        0d00h00m     ttyS0           5713  Local  network-engineer
130 vty 0    [C]test        0d00h01m     pts/0           5688  Local  network-engineer
OcNOS#
```

CLI Commands

aaa authentication login default fallback error

Use this command to enable fallback to local authentication for the default login if remote authentication is configured and all AAA servers are unreachable.

Use the `no` form of this command to disable fallback to local authentication.

Note: If you have specified `local` (use local authentication) in the aaa authentication login default command, you do not need to use this command to ensure that “fall back to local” occurs.

Command Syntax

```
aaa authentication login default fallback error local (vrf management|)
no aaa authentication login default fallback error local (vrf management|)
```

Parameters

`management`	Management VRF

Default

By default, AAA authentication is local.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#aaa authentication login default fallback error local vrf management
```

aaa authentication login default

Use this command to set the AAA authentication methods.

Use the `no` form of this command to set the default AAA authentication method (local).

Command Syntax

```
aaa authentication login default (vrf management|) ((group LINE) | (local (|none))
  | (none))
no aaa authentication login default (vrf management|) ((group) | (local (|none)) |
  (none))
```

Parameters

`group`		Use a server group list for authentication
	`LINE`	A space-separated list of up to 8 configured RADIUS or TACACS+, server group names followed by `local` or `none` or both `local` and `none`. The list can also include:
		`radius` All configured RADIUS servers
		`tacacs+` All configured TACACS+ servers
`local`		Use local authentication
`none`		No authentication
`management`		Management VRF

Default

By default, AAA authentication method is local

By default, groups: RADIUS or TACACS+

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3.

Examples

```
#configure terminal
(config)#aaa authentication login default vrf management group radius
```

Abbreviations

Acronym	Descripttions
AAA	accounting, authentication, authorization
RADIUS	Remote Authentication Dial-In User Service

Modified Extended ACL Deny Rule Behavior in VTY

Overview

The Access Control List refers to rules that allow or deny management protocols to control the network traffic, thus reducing network attacks from external sources.

Users can create Standard and Extended ACL rules and attach them to a virtual teletype (VTY) command line interface. These ACL rules are applied on both Management and Default virtual routing and forwarding (VRFs).

In the case of Standard ACLs, the permit/deny rules are applied only for management protocols such as Telnet/SSH/ SSH-Netconf protocols (port numbers 22,23,830).

Extended ACL rules are applied as configured by the user, and it is not limited to management protocols only, unlike Standard ACLs.

When a user configures a rule with 'deny any any any' and attaches it to the VTY, it effectively blocks only the Telnet, SSH, and NetConf protocols on the control plane

For example, when a user configures a rule as below and attach them to VTY, If the deny ACL rule includes ’any’ value in protocol, only Telnet/SSH/SSH-NetConf protocols are denied.

```
ip access-list ssh-access
10 permit tcp 10.12.43.0/24 any eq ssh
20 deny any any any
```

Note: To deny any protocols other than Telnet/SSH/SSH-Netconf, create a deny rule with the specific protocol access on VTY. For example: To deny OSPF protocol from all the source and destination address, apply the rule, `10 deny ospf any any`.

Feature Characteristics

In general, the VTY ACLs are more specific to management protocols. Hence, the Extended ACL “Any” rule translation is enhanced to allow management protocols as follows:

- If the **deny** ACL rule includes any value in protocol, only Telnet/SSH/SSH-Netconf protocols are denied.
- The **permit** ACL rule is unchanged.

Benefits

This feature allows the customer to define a Extended ACL deny rule only to the management protocol without impacting other control protocols.

Configure a separate Extended ACL deny rule to deny protocols other than Telnet, SSH, and NetConf.

Configuration

Refer to *Access Control Lists Configurations* section in the *System Management Configuration* guide, Release 6.4.1.

Implementation Examples

```
OcNOS#show running-config aclmgr
ip access-list ssh-access
 10 permit tcp 10.12.43.0/24 any eq ssh
 20 deny tcp 10.12.33.0/24 any eq 6513
 30 deny any 10.12.34.0/24 any
 40 deny any any any
!
line vty
 ip access-group ssh-access in

#####iptables o/p######

root@OcNOS:~# iptables -L
Chain INPUT (policy ACCEPT)
target     prot opt source               destination
ACCEPT     tcp  --  10.12.43.0/24        anywhere             tcp dpt:ssh
DROP       tcp  --  10.12.33.0/24        anywhere             tcp dpt:tls_netconf
DROP       tcp  --  10.12.34.0/24        anywhere             multiport dports
ssh,telnet,ssh_netconf
DROP       tcp  --  anywhere             anywhere             multiport dports
ssh,telnet,ssh_netconf
```

CLI Commands

Refer to *Access Control List Commands (Standard)* section of the *System Management Configuration* guide.

Abbreviations

Acronym	Expantion
ACL	Access control list
VRF	Virtual Routing Forwarding
VTY	Virtual teletype

Improved Management

Release 6.4.1

This section, describes the Network Monitoring, and DHCP group configurations introduced in the 6.4.1 release.

- Streaming Telemetry
- Support VLAN Range in SPAN
- Route Monitor
- DHCP Server Group

Streaming Telemetry

Overview

Streaming telemetry allows users to monitor network health by efficiently streaming operational data of interest from OcNOS routers. This structured data is transmitted to remote management systems for proactive network monitoring and understanding CPU and memory usage in managed devices for troubleshooting.

A machine learning (ML) database can be created with telemetry data to establish a baseline for normal network operation and predict or mitigate network issues.

Feature Characteristics

OcNOS version 6.4.1 introduces the initial features for Streaming Telemetry, which include support for gNMI-based Dial-in mode Telemetry for the management plane. The initial feature list includes support for the "**STREAM**" type and "**SAMPLING**" mode subscription for the Subscribe Remote Procedure Call (RPC). The gNMI-based collector connects to the OcNOS target device and invokes the Subscribe RPC, specifying the set of path(s) of interest. Below are the two key components involved:

- **gNMI Server (OcNOS Target):** The gNMI server operates within the OcNOS device, serving as the source of telemetry data. It supports the gNMI protocol, allowing gNMI-based clients (collectors) to request and receive streaming data. The server streams the requested data to the client according to the specified parameters.
- **gNMI Client (Collector):** The gNMI client, also known as the collector, runs outside the OcNOS target device and is responsible for receiving and gathering telemetry data. In this context, it is the entity that connects to the OcNOS target device to collect data using the gNMI protocol. The collector initiates the Subscribe RPC to specify the data of interest.

Figure 1 illustrates the gNMI client's (Collector) Subscribe request and response (RPC) interaction with the gNMI server (OcNOS Target).

Figure 1: Sample Subscribe Request

Dial-in Mode: Dial-in mode is the method used to establish a telemetry connection where the collector initiates the connection to the server. In this mode, the collector sends a Subscribe RPC request to the target device, and the server running on the target device streams the data to the collector.

Example Message Flow: Subscribe Request and Response

Figure 2 illustrates a sample gnmic Subscribe Request and Subscribe Response between the collector and the OcNOS target device.

Step 1: Subscription Request Initiation

- The gnmic collector server initiates a Subscribe Request by sending a Subscribe RPC in Stream type.
- This subscription request aims explicitly to gather data related to interface state counters and CPU state.
- A fixed 30/45-second sampling interval is set for data collection.

Step 2: Data Collection and Processing

- The gNMI server, within the OcNOS router, is responsible for data collection.
- At regular 30/45-second intervals, it retrieves data from the sensor path, focusing on interface state counters and CPU State.
- The received data undergoes a validation process, and the data is transformed into the required encoding type.

Step 3: Continuous Subscription Response Streaming

- The gNMI Server responds to the subscription request by continuously streaming Subscribe Response data.
- This streaming process maintains the same 30/45-second interval as the data collection.
- The collected data is streamed in real-time to the gnmic collector server.

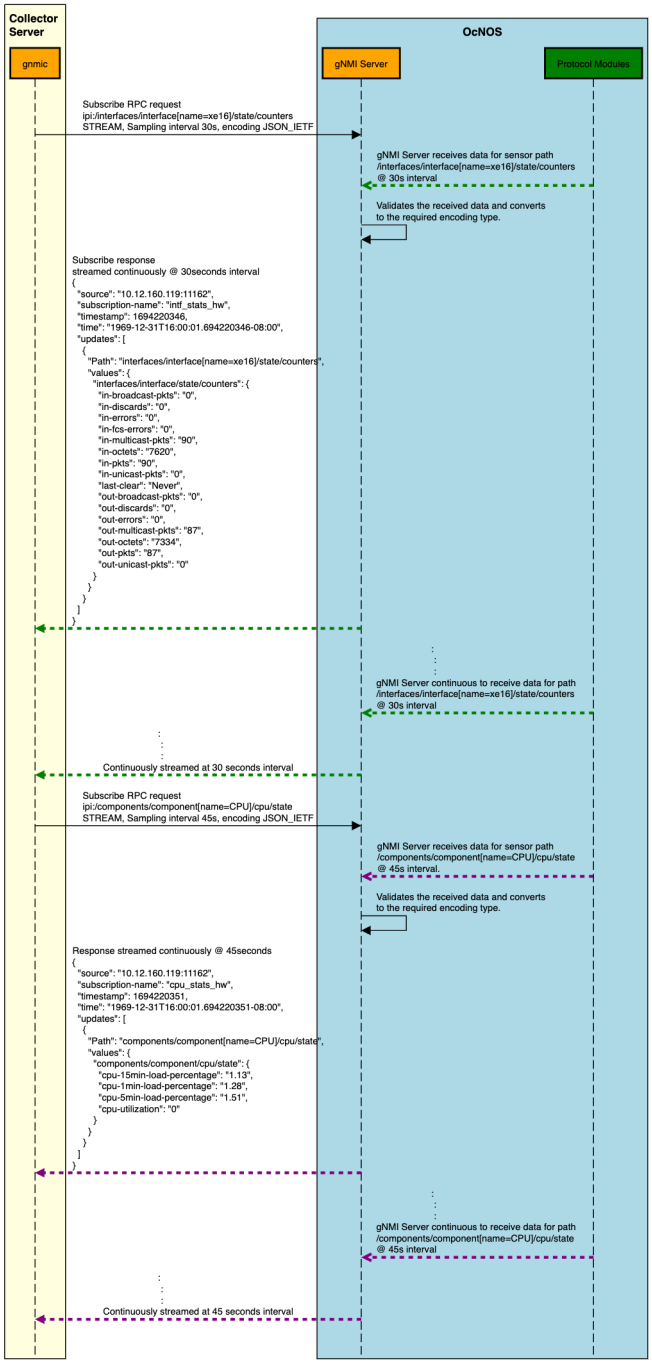

Figure 2: Message Flow: Subscribe Request and Response

Scale and Minimum Sample Interval Supported

To limit the impact of telemetry on critical features of the OcNOS target device, certain limits have been implemented. In Stream mode, there is a maximum limit of 100 sensor paths that can be subscribed to at any given point in time. Additionally, the minimum supported sample interval is 10 seconds.

Scale Scenarios

1. **New Subscribe RPC Request Makes Total Paths To Not Exceed 100:** When these new paths are added to the existing paths already handled by gNMI server, the total number does not exceed the maximum limit of 100 paths. Consequently, the gNMI server accepts this subscribe request and proceeds with the processing.
2. **New Subscribe RPC Request Makes Total Paths To Reach 100:** With the new Subscribe RPC Request, the total paths handled would be exactly equal to 100. The gNMI server accepts the new subscribe request; however, a warning is logged by the gNMI server, indicating that the maximum number of paths has been reached, and it signifies that no new Subscribe RPC Stream mode requests will be handled until the number of currently handled paths drops below 100.
3. **New Subscribe RPC Request Makes Total Paths To Exceed 100:** With the new Subscribe RPC Request, the total paths handled exceed 100. The gNMI server returns an error. The RPC request is not closed but will be accepted and responded to when the total number of paths handled drops to a level that can accommodate this RPC request.

Minimum Sample Interval: The minimum supported sample interval is 10 seconds. Any sampling mode request with a sample interval of less than 10 seconds will result in an error. However, if a sample interval is 0, it defaults to the minimum sample interval supported by the gNMI server, which is 10 seconds.

Benefits

Proactive Network Monitoring: Obtain real-time insights into network health and performance, and how to enable quicker response to issues.

Resource Utilization Monitoring: Monitor CPU and memory utilization to optimize resource allocation and performance.

Predictive Troubleshooting: Identify patterns and potential issues before they impact the network, reducing downtime.

Automation and Resilience: Use telemetry data to automate network management tasks and design a more resilient network.

Prerequisites

Before configuring Streaming Telemetry, ensure that:

- A supported OcNOS router running a compatible release.
- Access to the management interface of the router.
- Any gNMI client that complies with gNMI specifications can be used as a client.

Configuration

In this example, streaming telemetry with OcNOS is demonstrated, using '`gnmic`' as the gNMI Client.

gNMI Specification can be found at: https://github.com/openconfig/reference/blob/master/rpc/gnmi/gnmi-specification.md

The '`gnmic`' tool is available at: https://github.com/openconfig/gnmic

Figure 3: Streaming Telemetry Topology

gnmic installation

To install gnmic, use the following command:

```
bash -c "$(curl -sL https://get-gnmic.openconfig.net)"
```

To enable streaming telemetry on OcNOS:

```
OcNOS#configure terminal
OcNOS(config)#feature streaming-telemetry
OcNOS(config)#commit
```

Telemetry Subscription Request via gnmic Command and YAML Input

Use the gnmic command with a YAML file input to request telemetry subscriptions with multiple paths.

```
gnmic -a <ipaddress:port> -u <UserName> -p <Password> --insecure  --config <path to
  config file> subscribe
```

This command establishes a telemetry subscription with the specified paths defined in the YAML file.

Telemetry Subscription Request via gnmic Command with a Single Path Option

Use the gnmic command with a single path option to request a telemetry subscription for a specific data path.

```
gnmic -a  <ipaddress:port> -u <UserName> -p <Password> --encoding json_ietf --
  insecure --mode STREAM --stream-mode sample --sample-interval sample-interval-
  value sub --path <path>
```

This command creates a telemetry subscription for the specified path with the chosen sample interval and encoding format.

Supported gnmic Options

The below table explains the option fields.

gnmic Options details

Option	Description
--encoding	Specifies the encoding format (`JSON_IETF`).
--mode	Sets the mode of operation (`STREAM`).

Option	Description
--insecure	Allows insecure connections.
--stream-mode	Sets the stream mode (`Sample`).
--sample-interval	Sets the sample interval (`10s`). Note: Interval should be 10s or more.
--config	Specifies the YAML configuration file path (Example: `input_path.yaml`).
--path	Sets the path to subscribe to specific data (Example: `'ipi:/interfaces/interface[name=ce51]/state'`). Note: For multiple paths specify each path with --path option.
--prefix	Defines a common prefix for all specified paths (Example: `'ipi:/interfaces'`).

Invoking Subscribe RPC with gnmic

Use Case 1: Monitoring Interface State with Single Path Option

In this use case, gnmic subscribes to a specific path using the Subscribe RPC, monitoring the state of an interface with the path `'ipi:/interfaces/interface[name=ce51]/state'`.

```
#gnmic -a 10.12.91.111:11162 -u ocnos -p ocnos --encoding json_ietf --insecure
--mode STREAM --stream-mode sample --sample-interval 10s sub --path 'ipi:/
interfaces/interface[name=ce51]/state'

{
  "source": "10.12.91.111:11162",
  "subscription-name": "default-1695368813",
  "timestamp": 1551956933,
  "time": "1970-01-01T05:30:01.551956933+05:30",
  "updates": [
    {
      "Path": "ipi:interfaces/interface[name=ce51]/state",
      "values": {
        "interfaces/interface/state": {
          "admin-status": "up",
          "counters": {
            "in-broadcast-pkts": "0",
            "in-discards": "0",
            "in-errors": "0",
            "in-fcs-errors": "0",
            "in-multicast-pkts": "23",
            "in-octets": "2126",
            "in-pkts": "23",
            "in-unicast-pkts": "0",
            "last-clear": "Never",
            "out-broadcast-pkts": "0",
            "out-discards": "0",
            "out-errors": "0",
            "out-multicast-pkts": "28",
            "out-octets": "2552",
            "out-pkts": "28",
            "out-unicast-pkts": "0"
```

```
            },
            "ifindex": 10051,
            "last-change": 15500,
            "logical": false,
            "oper-status": "up"
          }
        }
      }
    ]
  }
```

The output of the Subscribe RPC includes the following information:

Subscribe RPC Output details

Option	Description
source	The source IP address and port of the gNMI server.
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 subscribed data.
values	The values of the subscribed data.

Validation

The below show command provides details about the subscriptions that have been established, including the client ID, sampling interval, encoding type, and the sensor paths that are being monitored.

```
OcNOS#show streaming-telemetry dynamic-subscriptions

 Feature streaming telemetry : Enabled

 SI:  Sampling Interval in seconds

 Enc-Type: Encoding type

 Dial-In Subscription Details:

 ClientIP:Port               ID     SI    Enc-Type       Origin:Path

 -------------            ------   ----   --------       ------------

 10.12.43.165:59304        4148   10     JSON_IETF       ipi:interfaces/interface[name=ce51]/state/counters

                                                          ipi:interfaces/interface[name=ce51]/state
```

Use Case 2: Monitoring Interface State with Multiple Path Option

In this use case, gnmic subscribes to a specific path using the Subscribe RPC, monitoring the state of an interface with the multiple path `'ipi:/interfaces/interface[name=ce51]/state'` and `'ipi:/interfaces/interface[name=ce52]/state'`.

```
#gnmic -a  10.12.91.111:11162 -u ocnos -p ocnos --encoding json_ietf  --
insecure   --mode STREAM --stream-mode sample --sample-interval 11s sub --path
'ipi:/interfaces/interface[name=ce51]/state' --path 'ipi:/interfaces/
interface[name=ce52]/state'
```

```
{
  "source": "10.12.91.111:11162",
  "subscription-name": "default-1695377304",
  "timestamp": 1551965423,
  "time": "1970-01-01T05:30:01.551965423+05:30",
  "updates": [
    {
      "Path": "ipi:interfaces/interface[name=ce51]/state",
      "values": {
        "interfaces/interface/state": {
          "admin-status": "up",
          "counters": {
            "in-broadcast-pkts": "0",
            "in-discards": "0",
            "in-errors": "0",
            "in-fcs-errors": "0",
            "in-multicast-pkts": "10",
            "in-octets": "1060",
            "in-pkts": "10",
            "in-unicast-pkts": "0",
            "last-clear": "Never",
            "out-broadcast-pkts": "0",
            "out-discards": "0",
            "out-errors": "0",
            "out-multicast-pkts": "10",
            "out-octets": "1020",
            "out-pkts": "10",
            "out-unicast-pkts": "0"
          },
          "ifindex": 10051,
          "last-change": 22500,
          "logical": false,
          "oper-status": "up"
        }
      }
    }
  ]
}

{
  "source": "10.12.91.111:11162",
  "subscription-name": "default-1695377304",
  "timestamp": 1551965423,
  "time": "1970-01-01T05:30:01.551965423+05:30",
  "updates": [
    {
      "Path": "ipi:interfaces/interface[name=ce52]/state",
      "values": {
        "interfaces/interface/state": {
          "admin-status": "up",
          "counters": {
            "in-broadcast-pkts": "0",
            "in-discards": "0",
            "in-errors": "0",
            "in-fcs-errors": "0",
            "in-multicast-pkts": "13",
```

```
                "in-octets": "1664",
                "in-pkts": "13",
                "in-unicast-pkts": "0",
                "last-clear": "Never",
                "out-broadcast-pkts": "0",
                "out-discards": "0",
                "out-errors": "0",
                "out-multicast-pkts": "10",
                "out-octets": "1020",
                "out-pkts": "10",
                "out-unicast-pkts": "0"
              },
              "ifindex": 10052,
              "last-change": 22500,
              "logical": false,
              "oper-status": "up"
            }
          }
        }
      ]
    }
```


Validation

The below show command provides details about the subscriptions that have been established, including the client ID, sampling interval, encoding type, and the sensor paths that are being monitored.

```
OcNOS#show streaming-telemetry dynamic-subscriptions

 Feature streaming telemetry : Enabled

 SI:  Sampling Interval in seconds

 Enc-Type: Encoding type

 Dial-In Subscription Details:

 ClientIP:Port               ID     SI    Enc-Type       Origin:Path

 -------------            ------   ----   --------       ------------

 10.12.43.145:59334        42000   11     JSON_IETF      ipi:interfaces/interface[name=ce52]/state/counters
                                                         ipi:interfaces/interface[name=ce52]/state
                                                         ipi:interfaces/interface[name=ce51]/state/counters
                                                         ipi:interfaces/interface[name=ce51]/state
```

YAML File Input for Multiple Path Subscription

Use Case 1: Configuring One Subscription Requests with Multiple Path Option

This use case illustrates the configuration of a subscription request with multiple paths using a YAML file input. It streamlines the subscription setup process by specifying the desired paths and subscription parameters directly in the YAML file.

YAML File Content (**single_request.yaml**)

```
#cat single_request.yaml
subscriptions:  # Container for subscriptions
  interface_stats_hw:  # A named subscription, where the key is the
subscription name
    paths:  # List of subscription paths for the named subscription
```

```
      - "ipi:/interfaces/interface[name=xe1]/state"
      - "ipi:/interfaces/interface[name=vlan1.10]/state"
    stream-mode: sample  # One of [on-change, target-defined, sample]
    sample-interval: 12s  # Sampling interval (e.g., 12 seconds)
    encoding: json_ietf  # Encoding format for telemetry data (e.g.,
JSON_IETF)
```

gnmic Command

```
# gnmic -a  10.12.91.111:11162 -u ocnos -p ocnos  --insecure  --config
single_request.yaml subscribe

{
  "source": "10.12.91.111:11162",
  "subscription-name": "interface_stats_hw",
  "timestamp": 1551965792,
  "time": "1970-01-01T05:30:01.551965792+05:30",
  "updates": [
    {
      "Path": "ipi:interfaces/interface[name=xe1]/state",
      "values": {
        "interfaces/interface/state": {
          "admin-status": "up",
          "counters": {
            "in-broadcast-pkts": "0",
            "in-discards": "0",
            "in-errors": "0",
            "in-fcs-errors": "0",
            "in-multicast-pkts": "0",
            "in-octets": "0",
            "in-pkts": "0",
            "in-unicast-pkts": "0",
            "last-clear": "Never",
            "out-broadcast-pkts": "0",
            "out-discards": "0",
            "out-errors": "0",
            "out-multicast-pkts": "2",
            "out-octets": "164",
            "out-pkts": "2",
            "out-unicast-pkts": "0"
          },
          "ifindex": 10001,
          "last-change": 0,
          "logical": false,
          "oper-status": "down"
        }
      }
    }
  ]
}

{
  "source": "10.12.91.111:11162",
  "subscription-name": "interface_stats_hw",
  "timestamp": 1551965792,
  "time": "1970-01-01T05:30:01.551965792+05:30",
  "updates": [
    {
```



```
          "Path": "ipi:interfaces/interface[name=vlan1.10]/state",
          "values": {
            "interfaces/interface/state": {
              "admin-status": "up",
              "counters": {
                "in-broadcast-pkts": "0",
                "in-discards": "0",
                "in-errors": "0",
                "in-fcs-errors": "0",
                "in-multicast-pkts": "0",
                "in-octets": "0",
                "in-pkts": "0",
                "in-unicast-pkts": "0",
                "last-clear": "Never",
                "out-broadcast-pkts": "0",
                "out-discards": "0",
                "out-errors": "0",
                "out-multicast-pkts": "0",
                "out-octets": "0",
                "out-pkts": "0",
                "out-unicast-pkts": "0"
              },
              "ifindex": 25010,
              "last-change": 22500,
              "logical": false,
              "oper-status": "up"
            }
          }
        }
      ]
    }
```


Validation

The below show command provides details about the subscriptions that have been established, including the client ID, sampling interval, encoding type, and the sensor paths that are being monitored.

```
OcNOS#show streaming-telemetry dynamic-subscriptions

 Feature streaming telemetry : Enabled

 SI:  Sampling Interval in seconds

 Enc-Type: Encoding type

 Dial-In Subscription Details:

 ClientIP:Port              ID     SI    Enc-Type       Origin:Path

 -------------            ------   ----   --------       ------------

 10.12.43.135:58208        45333   12     JSON_IETF      ipi:interfaces/interface[name=xe1]/state/counters
                                                         ipi:interfaces/interface[name=xe1]/state
                                                         ipi:interfaces/interface[name=vlan1.10]/state/counters
                                                         ipi:interfaces/interface[name=vlan1.10]/state
```

Use Case 2: Configuring Multiple Subscription Requests with Multiple Path Option

This use case illustrates the configuration of multiple subscription request with multiple paths using a YAML file input. It streamlines the subscription setup process by specifying the desired paths and subscription parameters directly in the YAML file.

YAML File Content (**multiple_subs.yaml**)

```
#cat multiple_subs.yaml
subscriptions:  # Container for subscriptions
  RAM_stats_hw:     # A named subscription for RAM statistics
    paths:      # List of subscription paths for the RAM_stats_hw subscription
      - "ipi:/components/component[name=RAM]/ram/state"
    stream-mode: sample  # Stream mode for RAM statistics
    sample-interval: 11s  # Sampling interval for RAM statistics (e.g., 11
seconds)
    encoding: json_ietf  # Encoding format for RAM statistics (e.g.,
JSON_IETF)

  storage_stats_hw:     # A named subscription for storage statistics
    paths:      # List of subscription paths for the storage_stats_hw
subscription
      - "ipi:/components/component[name=HARD-DISK]/storage/state"
    stream-mode: sample  # Stream mode for storage statistics
    sample-interval: 12s  # Sampling interval for storage statistics (e.g., 12
seconds)
    encoding: json_ietf  # Encoding format for storage statistics (e.g.,
JSON_IETF)

  power-supply_stats_hw:     # A named subscription for power supply
statistics
    paths:      # List of subscription paths for the power-supply_stats_hw
subscription
      - "ipi:/components/component[name=PSU-1]/power-supply/state"
      - "ipi:/components/component[name=PSU-2]/power-supply/state"
    stream-mode: sample  # Stream mode for power supply statistics
    sample-interval: 13s  # Sampling interval for power supply statistics
(e.g., 13 seconds)
    encoding: json_ietf  # Encoding format for power supply statistics (e.g.,
JSON_IETF)

  intf-tray_stats_hw:     # A named subscription for interface tray statistics
    paths:      # List of subscription paths for the intf-tray_stats_hw
subscription
      - "ipi:/interfaces/interface[name=xe1]/state"
      - "ipi:/interfaces/interface[name=vlan1.8]/state"
    stream-mode: sample  # Stream mode for interface tray statistics
    sample-interval: 14s  # Sampling interval for interface tray statistics
(e.g., 14 seconds)
    encoding: json_ietf  # Encoding format for interface tray statistics
(e.g., JSON_IETF)
```

gnmic Command

```
# gnmic -a  10.12.91.111:11162 -u ocnos -p ocnos  --insecure  --config
multiple_subs.yaml subscribe

{
  "source": "10.12.91.111:11162",
  "subscription-name": "ram_stats_hw",
  "timestamp": 1551967101,
  "time": "1970-01-01T05:30:01.551967101+05:30",
  "updates": [
    {
      "Path": "ipi:components/component[name=RAM]/ram/state",
      "values": {
```



```
        "components/component/ram/state": {
          "available-high-memory": "0",
          "available-memory": "14743",
          "buffers": "15",
          "current-process-count": 232,
          "free-swap": "0",
          "shared-memory": "8",
          "total-high-memory": "0",
          "total-memory": "16012",
          "total-swap": "0",
          "used-memory": "1269"
        }
      }
    }
  ]
}

{
  "source": "10.12.91.111:11162",
  "subscription-name": "storage_stats_hw",
  "timestamp": 1551967102,
  "time": "1970-01-01T05:30:01.551967102+05:30",
  "updates": [
    {
      "Path": "ipi:components/component[name=HARD-DISK]/storage/state",
      "values": {
        "components/component/storage/state": {
          "free-memory": "16908",
          "total-memory": "30208",
          "used-memory": "5020"
        }
      }
    }
  ]
}

{
  "source": "10.12.91.111:11162",
  "subscription-name": "power-supply_stats_hw",
  "timestamp": 1551967103,
  "time": "1970-01-01T05:30:01.551967103+05:30",
  "updates": [
    {
      "Path": "ipi:components/component[name=PSU-1]/power-supply/state",
      "values": {
        "components/component/power-supply/state": {
          "capacity": "650",
          "fan1-rpm": 24288,
          "operational-status": "not-present",
          "output-current": "8.28",
          "output-voltage": "12.07",
          "power-consumption": "99",
          "temperature-sensor1": "22",
          "temperature-sensor2": "28",
          "temperature-sensor3": "24"
        }
      }
```

```
      }
    ]
  }

  {
    "source": "10.12.91.111:11162",
    "subscription-name": "power-supply_stats_hw",
    "timestamp": 1551967103,
    "time": "1970-01-01T05:30:01.551967103+05:30",
    "updates": [
      {
        "Path": "ipi:components/component[name=PSU-2]/power-supply/state",
        "values": {
          "components/component/power-supply/state": {
            "operational-status": "running",
            "temperature-sensor1": "0",
            "temperature-sensor2": "0",
            "temperature-sensor3": "0"
          }
        }
      }
    ]
  }

  {
    "source": "10.12.91.111:11162",
    "subscription-name": "intf-tray_stats_hw",
    "timestamp": 1551967104,
    "time": "1970-01-01T05:30:01.551967104+05:30",
    "updates": [
      {
        "Path": "ipi:interfaces/interface[name=xe1]/state",
        "values": {
          "interfaces/interface/state": {
            "admin-status": "up",
            "counters": {
              "in-broadcast-pkts": "0",
              "in-discards": "0",
              "in-errors": "0",
              "in-fcs-errors": "0",
              "in-multicast-pkts": "0",
              "in-octets": "0",
              "in-pkts": "0",
              "in-unicast-pkts": "0",
              "last-clear": "Never",
              "out-broadcast-pkts": "0",
              "out-discards": "0",
              "out-errors": "0",
              "out-multicast-pkts": "5",
              "out-octets": "410",
              "out-pkts": "5",
              "out-unicast-pkts": "0"
            },
            "ifindex": 10001,
            "last-change": 0,
            "logical": false,
            "oper-status": "down"
```

```
            }
          }
        }
      ]
    }

    {
      "source": "10.12.91.111:11162",
      "subscription-name": "intf-tray_stats_hw",
      "timestamp": 1551967104,
      "time": "1970-01-01T05:30:01.551967104+05:30",
      "updates": [
        {
          "Path": "ipi:interfaces/interface[name=vlan1.8]/state",
          "values": {
            "interfaces/interface/state": {
              "admin-status": "up",
              "counters": {
                "in-broadcast-pkts": "0",
                "in-discards": "0",
                "in-errors": "0",
                "in-fcs-errors": "0",
                "in-multicast-pkts": "0",
                "in-octets": "0",
                "in-pkts": "0",
                "in-unicast-pkts": "0",
                "last-clear": "Never",
                "out-broadcast-pkts": "0",
                "out-discards": "0",
                "out-errors": "0",
                "out-multicast-pkts": "0",
                "out-octets": "0",
                "out-pkts": "0",
                "out-unicast-pkts": "0"
              },
              "ifindex": 25008,
              "last-change": 22500,
              "logical": false,
              "oper-status": "up"
            }
          }
        }
      ]
    }
```


Validation

The below show command provides details about the subscriptions that have been established, including the client ID, sampling interval, encoding type, and the sensor paths that are being monitored.

```
OcNOS#show streaming-telemetry dynamic-subscriptions

 Feature streaming telemetry : Enabled

 SI:  Sampling Interval in seconds

 Enc-Type: Encoding type

 Dial-In Subscription Details:
```

```
ClientIP:Port             ID      SI    Enc-Type       Origin:Path

-------------           ------   ----   --------       ------------

10.12.43.155:58267        9453     14    JSON_IETF      ipi:interfaces/interface[name=xe1]/state/counters
                                                        ipi:interfaces/interface[name=xe1]/state
                                                        ipi:interfaces/interface[name=vlan1.8]/state/counters
                                                        ipi:interfaces/interface[name=vlan1.8]/state
10.12.43.155:58114        31533    11    JSON_IETF      ipi:components/component[name=RAM]/ram/state
10.12.43.155:58345        3374     12    JSON_IETF      ipi:components/component[name=HARD-DISK]/storage/state
10.12.43.155:58222        35994    13    JSON_IETF      ipi:components/component[name=PSU-1]/power-supply/state
                                                        ipi:components/component[name=PSU-2]/power-supply/state
```

Use Case 3: Configuring Multiple Subscription Requests with Prefix Option

This use case illustrates the configuration of multiple subscription request with prefix option using a YAML file input. It streamlines the subscription setup process by specifying the desired paths and subscription parameters directly in the YAML file.

YAML File Content (**prefix_path.yaml**)

```
#cat prefix_path.yaml
subscriptions:  # Container for subscriptions
  RAM_stats_hw:     # A named subscription for RAM statistics
    prefix: "ipi:"  # Common prefix for paths in this subscription
    paths:      # List of subscription paths for the RAM_stats_hw subscription
      - "/components/component[name=RAM]/ram/state"
    stream-mode: sample  # Stream mode for RAM statistics
    sample-interval: 11s  # Sampling interval for RAM statistics (e.g., 11
seconds)
    encoding: json_ietf  # Encoding format for RAM statistics (e.g.,
JSON_IETF)

  intf-tray_stats_hw:     # A named subscription for interface tray statistics
    prefix: "ipi:"  # Common prefix for paths in this subscription
    paths:      # List of subscription paths for the intf-tray_stats_hw
subscription
      - "/interfaces/interface[name=xe1]/state"
      - "/interfaces/interface[name=vlan1.8]/state"
    stream-mode: sample  # Stream mode for interface tray statistics
    sample-interval: 14s  # Sampling interval for interface tray statistics
(e.g., 14 seconds)
    encoding: json_ietf  # Encoding format for interface tray statistics
(e.g., JSON_IETF)
```

gnmic Command

```
# gnmic -a  10.12.91.111:11162 -u ocnos -p ocnos  --insecure  --config
prefix_path.yaml subscribe
{
  "source": "10.12.91.111:11162",
  "subscription-name": "ram_stats_hw",
  "timestamp": 1551968637,
  "time": "1970-01-01T05:30:01.551968637+05:30",
  "updates": [
    {
      "Path": "components/component[name=RAM]/ram/state",
      "values": {
        "components/component/ram/state": {
          "available-high-memory": "0",
          "available-memory": "14793",
          "buffers": "16",
```



```
            "current-process-count": 231,
            "free-swap": "0",
            "shared-memory": "8",
            "total-high-memory": "0",
            "total-memory": "16012",
            "total-swap": "0",
            "used-memory": "1219"
          }
        }
      }
    ]
}

{
  "source": "10.12.91.111:11162",
  "subscription-name": "intf-tray_stats_hw",
  "timestamp": 1551968640,
  "time": "1970-01-01T05:30:01.55196864+05:30",
  "updates": [
    {
      "Path": "interfaces/interface[name=xe1]/state",
      "values": {
        "interfaces/interface/state": {
          "admin-status": "up",
          "counters": {
            "in-broadcast-pkts": "0",
            "in-discards": "0",
            "in-errors": "0",
            "in-fcs-errors": "0",
            "in-multicast-pkts": "0",
            "in-octets": "0",
            "in-pkts": "0",
            "in-unicast-pkts": "0",
            "last-clear": "Never",
            "out-broadcast-pkts": "0",
            "out-discards": "0",
            "out-errors": "0",
            "out-multicast-pkts": "9",
            "out-octets": "738",
            "out-pkts": "9",
            "out-unicast-pkts": "0"
          },
          "ifindex": 10001,
          "last-change": 0,
          "logical": false,
          "oper-status": "down"
        }
      }
    }
  ]
}

{
  "source": "10.12.91.111:11162",
  "subscription-name": "intf-tray_stats_hw",
  "timestamp": 1551968640,
  "time": "1970-01-01T05:30:01.55196864+05:30",
```

```
      "updates": [
        {
          "Path": "interfaces/interface[name=vlan1.8]/state",
          "values": {
            "interfaces/interface/state": {
              "admin-status": "up",
              "counters": {
                "in-broadcast-pkts": "0",
                "in-discards": "0",
                "in-errors": "0",
                "in-fcs-errors": "0",
                "in-multicast-pkts": "0",
                "in-octets": "0",
                "in-pkts": "0",
                "in-unicast-pkts": "0",
                "last-clear": "Never",
                "out-broadcast-pkts": "0",
                "out-discards": "0",
                "out-errors": "0",
                "out-multicast-pkts": "0",
                "out-octets": "0",
                "out-pkts": "0",
                "out-unicast-pkts": "0"
              },
              "ifindex": 25008,
              "last-change": 22500,
              "logical": false,
              "oper-status": "up"
            }
          }
        }
      ]
    }
```


Validation

The below show command provides details about the subscriptions that have been established, including the client ID, sampling interval, encoding type, and the sensor paths that are being monitored.

```
OcNOS#show streaming-telemetry dynamic-subscriptions

 Feature streaming telemetry : Enabled

 SI:  Sampling Interval in seconds

 Enc-Type: Encoding type

 Dial-In Subscription Details:

 ClientIP:Port              ID      SI    Enc-Type       Origin:Path

 -------------            ------   ----   --------       ------------

 10.12.43.154:50167        32137    11    JSON_IETF      ipi:components/component[name=RAM]/ram/state

 10.12.43.154:50614        36412    14    JSON_IETF     ipi:interfaces/interface[name=vlan1.8]/state/counters
                                              ipi:interfaces/interface[name=vlan1.8]/state
                                              ipi:interfaces/interface[name=xe1]/state/counters
                                              ipi:interfaces/interface[name=xe1]/state
```

Supported Datamodel and Sensor Paths

Streaming telemetry incrementally supports all IPI datamodels, with OcNOS version 6.4.1 introducing support for two IPI datamodels listed below. Telemetry supports only operational containers and a subset of leaf attributes. The Pyang tree output below illustrates the supported containers or leaves, along with a list of supported container-level paths.

ipi-platform

```
 +--rw components
     +--ro component* [name]
        +--ro name            -> ../state/name
        +--ro state
        |  +--ro name?                            string
        |  +--ro type?                            ipi-platform-
types:cmm_component_type_t
        |  +--ro location?                        string
        |  +--ro mfg-name?                        string
        |  +--ro mfg-date?                        yang:date-and-time
        |  +--ro description?                     string
        |  +--ro hardware-version?                string
        |  +--ro firmware-version?                string
        |  +--ro software-version?                string
        |  +--ro serial-no?                       string
        |  +--ro part-no?                         string
        |  +--ro removable?                       boolean
        |  +--ro oper-status?                     ipi-platform-
types:cmm_component_oper_status_t
        |  +--ro product-name?                    string
        |  +--ro asset-tag?                       string
        |  +--ro component-additional-details*   string
        |  +--ro parent?                          -> /components/component/
state/name
        |  +--ro empty?                           boolean
        |  +--ro memory
        |  |  +--ro available?   uint64
        |  |  +--ro utilized?    uint64
        |  +--ro board-fru
        |  |  +--ro board-name?        string
        |  |  +--ro board-serial-no?   string
        |  |  +--ro board-mfg-name?    string
        |  |  +--ro board-mfg-date?    yang:date-and-time
        |  +--ro temperature
        |     +--ro instant?                          decimal64
        |     +--ro min?                              decimal64
        |     +--ro max?                              decimal64
        |     +--ro avg?                              decimal64
        |     +--ro interval?                         uint32
        |     +--ro sensor-name?                      string
        |     +--ro sensor-index?                     uint8
        |     +--ro alarm-status?                     boolean
        |     +--ro alarm-threshold?                  decimal64
        |     +--ro alarm-severity?                   cml_alarm_severity_t
        |     +--ro minimum-emergency-temperature?    decimal64
        |     +--ro maximum-emergency-temperature?    decimal64
        |     +--ro minimum-alert-temperature?        decimal64
        |     +--ro maximum-alert-temperature?        decimal64
        |     +--ro minimum-critical-temperature?     decimal64
```

```
|     +--ro maximum-critical-temperature?    decimal64
+--ro cpu
|  +--ro state
|     +--ro cpu-1min-load-percentage?    decimal64
|     +--ro cpu-5min-load-percentage?    decimal64
|     +--ro cpu-15min-load-percentage?   decimal64
|     +--ro cpu-utilization?             decimal64
+--ro storage
|  +--ro state
|     +--ro total-memory?   uint64
|     +--ro used-memory?    uint64
|     +--ro free-memory?    uint64
+--ro ram
|  +--ro state
|     +--ro total-memory?            uint64
|     +--ro used-memory?             uint64
|     +--ro available-memory?        uint64
|     +--ro shared-memory?           uint64
|     +--ro buffers?                 uint64
|     +--ro total-swap?              uint64
|     +--ro free-swap?               uint64
|     +--ro current-process-count?   uint16
|     +--ro total-high-memory?       uint64
|     +--ro available-high-memory?   uint64
+--ro transceiver
|  +--ro state
|     +--ro grid-spacing?              decimal64
|     +--ro first-frequency?           decimal64
|     +--ro last-frequency?            decimal64
|     +--ro transceiver-temperature?   decimal64
|     +--ro transceiver-voltage?       decimal64
+--ro power-supply
|  +--ro state
|     +--ro operational-status?    cml_cmm_power_supply_operstatus_t
|     +--ro capacity?              decimal64
|     +--ro power-consumption?     decimal64
|     +--ro input-power?           decimal64
|     +--ro input-voltage?         decimal64
|     +--ro output-voltage?        decimal64
|     +--ro input-current?         decimal64
|     +--ro output-current?        decimal64
|     +--ro temperature-sensor1?   decimal64
|     +--ro temperature-sensor2?   decimal64
|     +--ro temperature-sensor3?   decimal64
|     +--ro fan1-rpm?              uint32
|     +--ro fan2-rpm?              uint32
|     +--ro fan3-rpm?              uint32
|     +--ro fan4-rpm?              uint32
+--ro fan
|  +--ro state
|     +--ro rpm?            uint32
|     +--ro fan-status?     cml_cmm_fan_status_t
|     +--ro fan-location?   cml_cmm_fan_location_t
+--ro fan-tray
   +--ro state
      +--ro status?   cml_cmm_fan_tray_status_t
```

ipi-interface

```
      +--rw interfaces
         +--rw interface* [name]
            +--rw name      -> ../config/name
            +--rw config
            |  +--rw name?   string
            +--ro state
               +--ro ifindex?        uint32
               +--ro admin-status?   ipi-if-types:if_interface_admin_status_t
               +--ro oper-status?    ipi-if-types:if_interface_oper_status_t
               +--ro last-change?    yang:timeticks
               +--ro logical?        boolean
               +--ro counters
                  +--ro in-octets?            yang:counter64
                  +--ro in-pkts?              yang:counter64
                  +--ro in-unicast-pkts?      yang:counter64
                  +--ro in-broadcast-pkts?    yang:counter64
                  +--ro in-multicast-pkts?    yang:counter64
                  +--ro in-discards?          yang:counter64
                  +--ro in-errors?            yang:counter64
                  +--ro in-fcs-errors?        yang:counter64
                  +--ro out-octets?           yang:counter64
                  +--ro out-pkts?             yang:counter64
                  +--ro out-unicast-pkts?     yang:counter64
                  +--ro out-broadcast-pkts?   yang:counter64
                  +--ro out-multicast-pkts?   yang:counter64
                  +--ro out-discards?         yang:counter64
                  +--ro out-errors?           yang:counter64
                  +--ro last-clear?           ipi-if-types:if_last_clear_time_t
```

Container Level Sensor Paths and Leaf Attributes

The below section lists the container level sensor paths and leaf attributes supported for telemetry.

ipi-interface

Interface State

```
    Sensor Path
            ipi:/interfaces/interface[name]/state

                /interfaces/interface[name]/state/name
                /interfaces/interface[name]/state/ifindex
                /interfaces/interface[name]/state/admin-status
                /interfaces/interface[name]/state/oper-status
                /interfaces/interface[name]/state/last-change
                /interfaces/interface[name]/state/logical
```

Interface Counters

```
    Sensor Path
            ipi:/interfaces/interface[name]/state/counters

                /interfaces/interface[name]/state/counters/in-octets
                /interfaces/interface[name]/state/counters/in-pkts
                /interfaces/interface[name]/state/counters/in-unicast-pkts
                /interfaces/interface[name]/state/counters/in-broadcast-pkts
                /interfaces/interface[name]/state/counters/in-multicast-pkts
                /interfaces/interface[name]/state/counters/in-discards
```

```
/interfaces/interface[name]/state/counters/in-errors
/interfaces/interface[name]/state/counters/in-fcs-errors
/interfaces/interface[name]/state/counters/out-octets
/interfaces/interface[name]/state/counters/out-pkts
/interfaces/interface[name]/state/counters/out-unicast-pkts
/interfaces/interface[name]/state/counters/out-broadcast-pkts
/interfaces/interface[name]/state/counters/out-multicast-pkts
/interfaces/interface[name]/state/counters/out-discards
/interfaces/interface[name]/state/counters/out-errors
/interfaces/interface[name]/state/counters/last-clear
```

ipi-platform

The paths listed below represent telemetry paths for monitoring the state of various components, including CPU, storage, RAM, power supply, fans, fan trays, and transceivers.

CPU

```
Sensor Path
        ipi:/components/component[name]/cpu/state

Leaf Attributes
            /components/component[name]/cpu/state/cpu-1min-load-percentage
            /components/component[name]/cpu/state/cpu-5min-load-percentage
            /components/component[name]/cpu/state/cpu-15min-load-percentage
            /components/component[name]/cpu/state/cpu-utilization
```

Storage

```
Sensor Path
        ipi:/components/component[name]/storage/state/

Leaf Attributes
            /components/component[name]/storage/state/total-memory
            /components/component[name]/storage/state/used-memory
            /components/component[name]/storage/state/free-memory
```

RAM

```
Sensor Path
        ipi:/components/component[name]/ram/state/

Leaf Attributes
            /components/component[name]/ram/state/total-memory
            /components/component[name]/ram/state/used-memory
            /components/component[name]/ram/state/available-memory
            /components/component[name]/ram/state/shared-memory
            /components/component[name]/ram/state/buffers
            /components/component[name]/ram/state/total-swap
            /components/component[name]/ram/state/free-swap
            /components/component[name]/ram/state/current-process-count
            /components/component[name]/ram/state/total-high-memory
            /components/component[name]/ram/state/available-high-memory
```

Power-Supply

```
Sensor Path
        ipi:/components/component[name]/power-supply/state/

Leaf Attributes
```

```
            /components/component[name]/power-supply/state/capacity
            /components/component[name]/power-supply/state/power-consumption
            /components/component[name]/power-supply/state/input-power
            /components/component[name]/power-supply/state/input-voltage
            /components/component[name]/power-supply/state/input-current
            /components/component[name]/power-supply/state/output-voltage
            /components/component[name]/power-supply/state/output-current
            /components/component[name]/power-supply/state/operational-status
            /components/component[name]/power-supply/state/fan1-rpm
            /components/component[name]/power-supply/state/fan2-rpm
            /components/component[name]/power-supply/state/fan3-rpm
            /components/component[name]/power-supply/state/fan4-rpm
            /components/component[name]/power-supply/state/temperature-sensor1
            /components/component[name]/power-supply/state/temperature-sensor2
            /components/component[name]/power-supply/state/temperature-sensor3
```

Fan

```
    Sensor Path
            ipi:/components/component[name]/fan/state/

    Leaf Attributes
            /components/component[name]/fan/state/rpm
            /components/component[name]/fan/state/fan-status
            /components/component[name]/fan/state/fan-location
```

Fan-Tray

```
    Sensor Path
            ipi:/components/component[name]/fan-tray/state/

    Leaf Attributes
            /components/component[name]/fan-tray/state/status
```

Transceiver

```
    Sensor Path
            ipi:/components/component[name]/transceiver/state/

    Leaf Attributes
            /components/component[name]/transceiver/state/grid-spacing
            /components/component[name]/transceiver/state/first-frequency
            /components/component[name]/transceiver/state/last-frequency
            /components/component[name]/transceiver/state/transceiver-
    temperature
            /components/component[name]/transceiver/state/transceiver-voltage
```

Platform State

```
    Sensor Path
            ipi:/components/component[name]/state/

    Leaf Attributes
            /components/component[name]/state/name
            /components/component[name]/state/type
            /components/component[name]/state/location
            /components/component[name]/state/mfg-name
            /components/component[name]/state/description
            /components/component[name]/state/hardware-version
            /components/component[name]/state/firmware-version
```

```
            /components/component[name]/state/software-version
            /components/component[name]/state/serial-no
            /components/component[name]/state/part-no
            /components/component[name]/state/removable
            /components/component[name]/state/oper-status
            /components/component[name]/state/product-name
            /components/component[name]/state/asset-tag
            /components/component[name]/state/component-additional-details
            /components/component[name]/state/parent
            /components/component[name]/state/empty

Sensor Path
        ipi:/components/component[name]/state/memory
Leaf Attributes
            /components/component[name]/state/memory/available
            /components/component[name]/state/memory/utilized

Sensor Path
        ipi:/components/component[name]/state/board-fru
Leaf Attributes
            /components/component[name]/state/board-fru/board-name
            /components/component[name]/state/board-fru/board-serial-no
            /components/component[name]/state/board-fru/board-mfg-name
            /components/component[name]/state/board-fru/board-mfg-date


Sensor Path
        ipi:/components/component[name]/state/temperature
Leaf Attributes
            /components/component[name]/state/temperature/instant
            /components/component[name]/state/temperature/min
            /components/component[name]/state/temperature/max
            /components/component[name]/state/temperature/avg
            /components/component[name]/state/temperature/interval
            /components/component[name]/state/temperature/sensor-name
            /components/component[name]/state/temperature/sensor-index
            /components/component[name]/state/temperature/alarm-status
            /components/component[name]/state/temperature/alarm-threshold
            /components/component[name]/state/temperature/alarm-severity
            /components/component[name]/state/temperature/minimum-emergency-
temperature
            /components/component[name]/state/temperature/maximum-emergency-
temperature
            /components/component[name]/state/temperature/minimum-alert-
temperature
            /components/component[name]/state/temperature/maximum-alert-
temperature
            /components/component[name]/state/temperature/minimum-critical-
temperature
            /components/component[name]/state/temperature/maximum-critical-
temperature
```

Implementation Examples

Typical Use Cases

- Enable Streaming Telemetry to monitor interface counters and the health of the OcNOS target device, including memory, CPU usage, fan speed, and temperature.
- Use telemetry data to trigger automated network tasks based on specific conditions.

Integration with Existing Features

Streaming Telemetry can be used in conjunction with other network monitoring and management features.

New CLI Commands

The Streaming Telemetry introduces the following configuration commands.

debug cml

Use this command to enable or disable debugging information for CML streaming telemetry.

Command Syntax

```
debug cml enable telemetry
debug cml disable telemetry
```

Parameters

None

Default

By default, debugging information is disabled.

Command Mode

Exec Mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The following example illustrates how to enable and disable the telemetry debugging information.

```
OcNOS#debug cml enable telemetry
OcNOS#debug cml disable telemetry
```

debug telemetry gnmi

Use this command to enable or disable gNMI server debugging logs with severity levels.

Command Syntax

```
debug telemetry gnmi (enable) (severity (debug|info|warning|error|fatal|panic|d-panic)|)
debug telemetry gnmi (disable) (severity (debug|info|warning|error|fatal|panic|d-panic)|)
```

Parameters

`debug`	Logs a message at debug level
`info`	Logs a message at info level
`warning`	Logs a message at warning level
`error`	Logs a message at error level
`fatal`	Logs a message and causes the program to exit with return code 1.
`panic`	Logs a message and triggers the program to generate a traceback.
`d-panic`	Logs at the Panic level

Default

By default, this command is disabled, and the gNMI server debugging level in the disabled state is set to the Error level.

Command Mode

Configure Mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The following example illustrates how to enable and disable the telemetry debug logs and their corresponding show output.

```
OcNOS(config)#feature streaming-telemetry
OcNOS(config)#debug telemetry gnmi enable severity warning
OcNOS(config)#commit
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
debug telemetry gnmi enable severity warning
!
OcNOS(config)#debug telemetry gnmi disable severity warning
OcNOS(config)#commit
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
!
```

feature streaming-telemetry

Use this command to enable the streaming telemetry and, upon configuration, to start the gNMI server. The gNMI server initiates listening for incoming gRPC connections on port 11162.

Use the no parameter of this command to disable the streaming telemetry, It will stop the gNMI server.

Command Syntax

```
feature streaming-telemetry
no feature streaming-telemetry
```

Parameters

None

Default

By default, the streaming-telemetry feature is disabled.

Command Mode

Configure mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The following example illustrates how to enable the streaming telemetry.

```
OcNOS#configure terminal
OcNOS(config)#feature streaming-telemetry
OcNOS(config)#commit
```

show running-config streaming-telemetry

Use this command to display streaming telemetry status in the running configuration.

Command Syntax

```
show running-config streaming-telemetry
```

Parameters

None

Command Mode

Exec mode and Configuration Mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The following example shows the streaming telemetry status in the `show running-config` output.

```
OcNOS#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#feature streaming-telemetry
OcNOS(config)#commit
OcNOS(config)#show running-config streaming-telemetry
!
feature streaming-telemetry
```

```
!
OcNOS(config)#exit
OcNOS#show running-config streaming-telemetry
!
feature streaming-telemetry
!
```

show streaming-telemetry dynamic-subscriptions

Use this command to display the streaming telemetry dial-in configurations.

Command syntax

```
show streaming-telemetry dynamic-subscriptions
```

Parameters

None

Command Mode

Exec mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The following example displays the streaming telemetry dial-in configuration output.

```
OcNOS#show streaming-telemetry dynamic-subscriptions

 Feature streaming telemetry : Enabled

 SI:  Sampling Interval in seconds

 Enc-Type: Encoding type

 Dial-In Subscription Details:

 ClientIP:Port               ID     SI    Enc-Type       Origin:Path

 -------------            ------   ----   --------       ------------

 10.12.43.175:59108        12396   10     JSON_IETF      ipi:interfaces/interface[name=eth0]/state/counters

                                                         ipi:interfaces/interface[name=eth0]/state

 10.12.43.175:59114        6001    15     JSON_IETF      ipi:components/component[name=CPU]/cpu/state
```

The below table explains the output fields.

show streaming-telemetry dynamic-subscriptions parameters output details

Field	Description
Feature streaming telemetry	Marked as “Enabled” confirms that streaming telemetry is active on the device.
Dial-In Subscription Details	Check the Dial-in subscription details.

show streaming-telemetry dynamic-subscriptions parameters output details

Field	Description
ClientIP: Port	Verify that the client IP and port listed matches the client that should be receiving telemetry data.
SI: Sampling-interval	Confirm that the sampling interval matches the desired frequency at which data is collected and sent.
Enc-type: Encoding-type	Ensure that the encoding type (e.g., JSON_IETF) matches the expected format for telemetry data.
Origin:Path	Review the sensor paths to ensure that they correspond to the specific data sources or paths of interest.

Troubleshooting

Follow the below troubleshooting steps, to debug telemetry related issues:

Verify Collector (gnmic) Command Options: Verify the input parameters, such as the sensor path, prefix and origin "`ipi:`".

Check the Encoding Method Compatibility: Check that the request conforms to the supported JSON-IETF encoding method.

Ensure Proper Connectivity: Validate the connectivity between the router and the remote management system. This involves verifying network settings, ports, firewalls, and any potential disruptions in communication.

Collector: If `gnmic` does not receive a response or not receiving expected response, restart the request using the "`--log`" option. If more verbose debug output is needed, consider adding the "`--debug`" option as well. The gnmic tool displays the possible cause for any error, which helps in debugging the issue.

gNMI Server: If the issue is on server side, follow the steps below to troubleshoot telemetry issues on the OcNOS target. Enable debug and verify the logs in `/var/log/messages` file.

1. In configure mode, enable debug with a specific severity level either "`info`" or "`debug`" level, using the following command:

   ```
   debug telemetry gnmi (enable) (severity
   (debug|info|warning|error|fatal|panic|d-panic)|)
   ```

Note: To disable the debug telemetry, configure `debug telemetry gnmi (disable)` command.

2. In Exec mode, enable telemetry related debugs, using the following command:

   ```
   debug cml enable telemetry
   ```

Note: To disable telemetry related debugs, configure "`debug cml disable telemetry`" command.

3. Collect the output of the following command to check the state of streaming telemetry:

   ```
   show streaming-telemetry dynamic dynamic-subscriptions
   ```

Note: If telemetry is in "`disabled`" state, then telemetry feature need to enabled.

4. Collect the output of the following command to gather diagnostic information and the logs in `/var/log/messages` file, to triage further.

   ```
   show techsupport all
   ```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
JSON	JavaScript Object Notation
RPC	Remote Procedure Call
gNMI	gRPC Network Management Interface
JSON-IETF	JSON-Internet Engineering Task Force

Glossary

The following provides definitions for key terms used throughout this document.

Streaming Telemetry	A monitoring approach that efficiently transmits operational data from OcNOS routers to remote management systems in real-time for analysis, troubleshooting, and network monitoring.
Telemetry Data	Structured operational data generated by routers that is transmitted in real-time to external systems for analysis.
JSON-IETF	JSON-IETF is a data interchange format that follows the specifications defined by the IETF. It is a lightweight, text-based format used for representing structured data. JSON-IETF is commonly used for configuration and data exchange in various network and Internet-related protocols.
Remote Management System	An external system responsible for monitoring, managing, and analyzing data received from network devices.
Network Health	The overall condition and performance of a network, including factors like stability, resource utilization, and data flow.
Resilient Network	A network designed to withstand failures or disruptions, maintaining functionality even in challenging conditions.

Support VLAN Range in SPAN

Overview

The Switch Port Analyzer (SPAN) monitors the traffic on source port and sends a copy of the traffic to a destination port. The network analyzer, which is attached to the destination port, analyzes the received traffic. The source port can either be a single port or multiple ports. A replication of the packets is sent to the destination port for analysis.

The SPAN is also referred to as port mirroring or port monitoring. It is installed in Layer 2 Access Control List (ACL) group by default. It is used for monitoring Ingress MAC ACL or VLAN group. Any packet received can be monitored based on source port including Physical or MAC or VLAN port.

This is an existing VLAN monitor session feature in the OcNOS DC, enhanced in current release to support VLAN ranges.

The following two CLIs are updated to support the VLAN ranges:

- hardware-profile filter (XGS)
- filter

Feature Characteristics

The VLAN range is supported only for ingress traffic.

LIMITATIONS

The ingress port mirroring is not supported for sub-interface and Switched Virtual Interface (SVI) interface.

Benefits

Users can apply port monitoring rules for multiple source ports, multiple VLANs, and a combination of port and VLAN ranges.

Configuration

To configure an ingress VLAN monitor session using VLAN ranges, perform the following configurations:

Topology

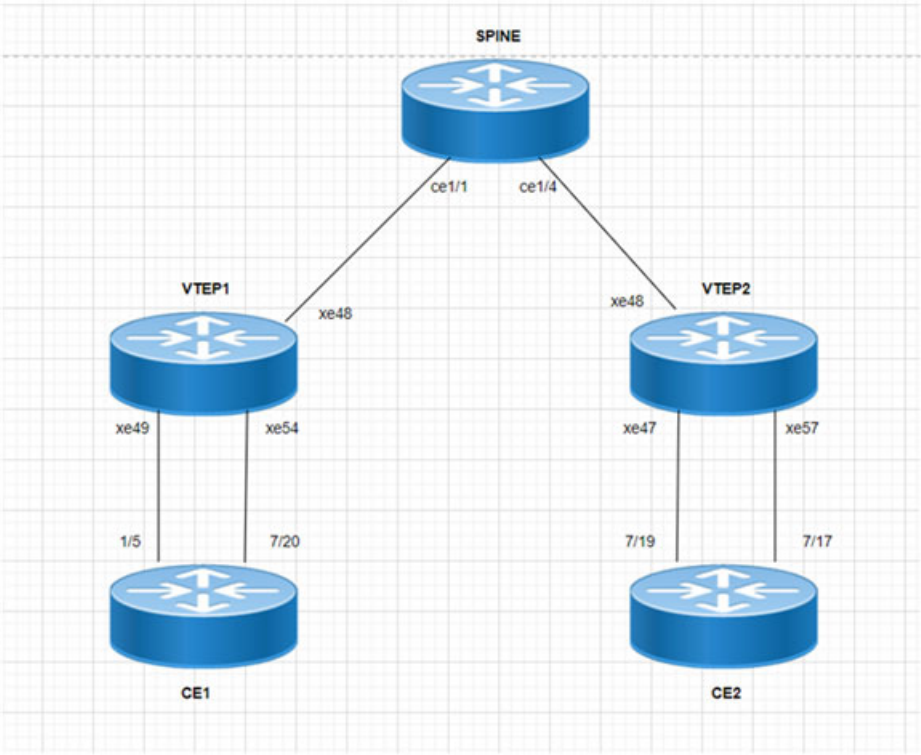

SPAN Topology

VTEP1

`VTEP1#configure terminal`	Enter configure mode.
`VTEP1(config)#hardware-profile filter ingress-mirror enable`	Enable hardware profile igress mirror.
`VTEP1(config)#nvo vxlan enable`	Enable vxlan.
`VTEP1(config)#evpn esi hold-time 60`	Configure esi hold timer.
`VTEP1(config)#evpn vxlan multihoming enable`	Enable VxLAN multihoming.
`VTEP1(config)#mac vrf VRF1`	Configure MAC VRF as VRF1.
`VTEP1(config-vrf)#rd 1.1.1.1:11`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 9.9.9.9:100`	Configure `route-target import and export.`
`VTEP1(config)#mac vrf VRF2`	Configure MAC VRF as VRF2.
`VTEP1(config-vrf)#rd 1.1.1.1:21`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 90.90.90.90:100`	Configure `route-target import and export.`

`VTEP1(config)#mac vrf VRF3`	Configure MAC VRF as VRF3.
`VTEP1(config-vrf)#rd 1.1.1.1:22`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 90.90.90.90:101`	Configure `route-target import and export.`
`VTEP1(config)#mac vrf VRF4`	Configure MAC VRF as VRF4.
`VTEP1(config-vrf)#rd 1.1.1.1:23`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 10.10.10.10:100`	Configure `route-target import and export.`
`VTEP1(config)#mac vrf VRF5`	Configure MAC VRF as VRF5.
`VTEP1(config-vrf)#rd 1.1.1.1:24`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 20.20.20.20:100`	Configure `route-target import and export.`
`VTEP1(config)#mac vrf VRF6`	Configure MAC VRF as VRF6.
`VTEP1(config-vrf)#rd 1.1.1.1:25`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 30.30.30.30:100`	Configure `route-target import and export.`
`VTEP1(config)#mac vrf VRF7`	Configure MAC VRF as VRF7.
`VTEP1(config-vrf)#rd 1.1.1.1:26`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 40.40.40.40:100`	Configure `route-target import and export.`
`VTEP1(config-vrf)#exit`	Exit from VRF mode
`VTEP1(config)#mac vrf VRF8`	Configure MAC VRF as VRF8
`VTEP1(config-vrf)#rd 1.1.1.1:27`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 50.50.50.50:100`	Configure `route-target import and export.`
`VTEP1(config-vrf)#exit`	Exit from VRF mode.
`VTEP1(config)#mac vrf VRF9`	Configure MAC VRF as VRF2.
`VTEP1(config-vrf)#rd 1.1.1.1:28`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 60.60.60.60:100`	Configure `route-target import and export.`
`VTEP1(config-vrf)#exit`	Exit from VRF mode.
`VTEP1(config)#mac vrf VRF10`	Configure MAC VRF as VRF2.
`VTEP1(config-vrf)#rd 1.1.1.1:29`	Configure `route distinguisher value.`
`VTEP1(config-vrf)#route-target both 70.70.70.70:100`	Configure `route-target import and export.`
`VTEP1(config-vrf)#exit`	Exit from VRF mode.
`VTEP1(config)#nvo vxlan vtep-ip-global 1.1.1.1`	Enable VxLAN Source VTEP IPp address global configuration.
`VTEP1(config)#nvo vxlan id 10 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.

`VTEP1(config)#nvo vxlan id 20 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF2`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP1(config)#nvo vxlan id 21 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF3`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP1(config)#nvo vxlan id 23 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF4`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP1(config)#nvo vxlan id 24 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF5`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP1(config)#nvo vxlan id 25 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#VxLAN  host-reachability-protocol evpn-bgp VRF6`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP1(config)#nvo VxLAN  id 26 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#VxLAN  host-reachability-protocol evpn-bgp VRF7`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP1(config)#nvo VxLAN  id 27 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF8`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP1(config)#nvo vxlan id 28 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF9`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP1(config)#nvo vxlan id 29 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP1(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF10`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP1(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP1(config)#qos enable`	Enable QoS.
`VTEP1(config)#hostname VTEP1`	Configure `system's network name as VTEP1`

VTEP1(config)#interface lo	Enter loopback interface mode.
VTEP1(config-if)#ip address 1.1.1.1/32 secondary	Configure the secondary IP address of the- loopback interface
VTEP1(config)#interface xe48	Enter interface mode.
VTEP1(config-if)#load-interval 30	Configure load interval.
VTEP1(config-if)#ip address 10.10.10.1/24	Configure the IP address of the interface.
VTEP1(config-if)#exit	Exit from interface mode.
VTEP1(config)#interface xe49	Enter interface mode.
VTEP1(config-if)#switchport	Enter the switchport mode.
VTEP1(config-if)#load-interval 30	Configure load interval.
VTEP1(config-if)#exit	Exit from interface mode.
VTEP1(config)#interface xe54	Enter interface mode.
VTEP1(config-if)#switchport	Enter the switchport mode.
VTEP1(config-if)#load-interval 30	Configure load interval.
VTEP1(config-if)#exit	Exit from interface mode.
VTEP1(config)#router ospf 100	Configure router ospf process ID.
VTEP1(config-router)#ospf router-id 1.1.1.1	Configure OSPF router id
VTEP1(config-router)#bfd all-interfaces	Enable BFD all interfaces
VTEP1(config-router)#network 1.1.1.1/32 area 0.0.0.0	Configure network and area as 0
VTEP1(config-router)#network 10.10.10.0/24 area 0.0.0.0	Configure network and area as 0
VTEP1(config-router)#exit	Exit from router ospf mode
VTEP1(config)#router bgp 500	Configure router bgp AS number
VTEP1(config-router)#bgp router-id 1.1.1.1	Configure BGP router ID.
VTEP1(config-router)#neighbor 2.2.2.2 remote-as 500	Configure a neighbor router and Peer AS Specify AS number of BGP neighbor.
VTEP1(config-router)#neighbor 2.2.2.2 update-source lo	Configure a neighbor router and Source of routing updates as loopbacl
VTEP1(config-router)#neighbor 2.2.2.2 advertisement-interval 0	Configure a neighbor router and minimum interval between sending BGP routing updates
VTEP1(config-router)#address-family ipv4 unicast	Enter Address Family command mode
VTEP1(config-router-af)#network 1.1.1.1/32	Configure a network to announce via BGP
VTEP1(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor
VTEP1(config-router-af)#exit-address-family	Exit from address family mode
VTEP1(config-router)#address-family l2vpn evpn	Enter Address Family with l2vpn evpn Identifier
VTEP1(config-router-af)#neighbor 2.2.2.2 activate	Activate the neighbor
VTEP1(config-router-af)#exit-address-family	Exit from address family mode
VTEP1(config-router)#exit	Exit from router bgp mode
VTEP1(config)#monitor session 1	Configure Ethernet SPAN session with preferences

VTEP1(config-monitor)#source interface xe49 rx	Configure source interface as Ingress
VTEP1(config-monitor)#destination interface xe54	Configure destination interface.
VTEP1(config-monitor)#10 filter vlan 2-6	Configure sequence number with filter option and specify the vlan ranges.
VTEP1(config-monitor)#no shut	Unshut a monitor session.
VTEP1(config-monitor)#exit	Exit from monitor session.
VTEP1(config)#nvo vxlan max-cache-disable 2500	Configure vxlan Max number of ARP/ND cache disable allowed for port-vlan
VTEP1(config)#nvo vxlan access-if port-vlan xe49 2	Configure VxLAN access-if single tagged interface name with VLAN id
VTEP1(config-nvo-acc-if)#map vnid 22	Map access port attribute with VxLAN Identifier.
VTEP1(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP1(config)#nvo vxlan access-if port-vlan xe49 3	Configure VxLAN access-if single tagged interface name with VLAN id
VTEP1(config-nvo-acc-if)#map vnid 23	Map access port attribute with VxLAN Identifier.
VTEP1(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP1(config)#nvo vxlan access-if port-vlan xe49 4	Configure VxLAN access-if single tagged interface name with VLAN id
VTEP1(config-nvo-acc-if)#map vnid 24	Map access port attribute with VxLAN Identifier.
VTEP1(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP1(config)#nvo vxlan access-if port-vlan xe49 5	Configure VxLAN access-if single tagged interface name with VLAN id
VTEP1(config-nvo-acc-if)#map vnid 25	Map access port attribute with VxLAN Identifier.
VTEP1(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP1(config)#nvo vxlan access-if port-vlan xe49 6	Configure VxLAN access-if single tagged interface name with VLAN id
VTEP1(config-nvo-acc-if)#map vnid 26	Map access port attribute with VxLAN Identifier.
VTEP1(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP1(config)#nvo vxlan access-if port-vlan xe49 7	Configure VxLAN access-if single tagged interface name with VLAN id
VTEP1(config-nvo-acc-if)#map vnid 27	Map access port attribute with VxLAN Identifier.
VTEP1(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP1(config)#nvo vxlan access-if port-vlan xe49 8	Configure VxLAN access-if single tagged interface name with VLAN id
VTEP1(config-nvo-acc-if)#map vnid 28	Map access port attribute with VxLAN Identifier.
VTEP1(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP1(config)#nvo vxlan access-if port-vlan xe49 9	Configure VxLAN access-if single tagged interface name with VLAN id
VTEP1(config-nvo-acc-if)#map vnid 29	Map access port attribute with VxLAN Identifier.
VTEP1(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP1(config)#nvo vxlan access-if port-vlan xe49 10	Configure VxLAN access-if single tagged interface name with VLAN id.

`VTEP1(config-nvo-acc-if)#map vnid 10`	Map access port attribute with VxLAN Identifier.
`VTEP1(config)#nvo vxlan access-if port-vlan xe49 11`	Configure VxLAN access-if `single tagged interface name with VLAN id.`
`VTEP1(config-nvo-acc-if)#map vnid 21`	Map access port attribute with VxLAN Identifier.
`VTEP1(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP1(config)#nvo vxlan access-if port-vlan xe49 12`	Configure VxLAN access-if `single tagged interface name with VLAN id.`
`VTEP1(config-nvo-acc-if)#map vnid 20`	Map access port attribute with VxLAN Identifier.
`VTEP1(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP1(config)#commit`	Commit the candidate configuration to the running configuration.

VTEP2

VTEP2#configure terminal	Enter configure mode.
VTEP2(config)#hardware-profile filter ingress-mirror enable	Enable hardware profile igress mirror
VTEP2(config)#nvo vxlan enable	Enable vxlan
VTEP2(config)#evpn esi hold-time 60	Config esi hold timer
VTEP2(config)#evpn vxlan multihoming enable	Enable vxlan multihoming
VTEP2(config)#mac vrf VRF1	Configure mac vrf as VRF1
VTEP2(config-vrf)#rd 2.2.2.2:11	Configure route distinguisher value.
VTEP2(config-vrf)#route-target both 9.9.9.9:100	Configure route-target import and export
VTEP2(config-vrf)#exit	Exit from vrf mode
VTEP2(config)#mac vrf VRF2	Configure mac vrf as VRF2
VTEP2(config-vrf)#rd 2.2.2.2:21	Configure route distinguisher value.
VTEP2(config-vrf)#route-target both 90.90.90.90:100	Configure route-target import and export
VTEP2(config-vrf)#exit	Exit from vrf mode
VTEP2(config)#mac vrf VRF3	Configure mac vrf as VRF3
VTEP2(config-vrf)#rd 2.2.2.2:22	Configure route distinguisher value.
VTEP2(config-vrf)#route-target both 90.90.90.90:101	Configure route-target import and export
VTEP2(config-vrf)#exit	Exit from vrf mode
VTEP2(config)#mac vrf VRF4	Configure mac vrf as VRF4
VTEP2(config-vrf)#rd 2.2.2.2:23	Configure route distinguisher value.
VTEP2(config-vrf)#route-target both 10.10.10.10:100	Configure route-target import and export
VTEP2(config-vrf)#exit	Exit from vrf mode
VTEP2(config)#mac vrf VRF5	Configure mac vrf as VRF5
VTEP2(config-vrf)#rd 2.2.2.2:24	Configure route distinguisher value.
VTEP2(config-vrf)#route-target both 20.20.20.20:100	Configure route-target import and export
VTEP2(config-vrf)#exit	Exit from vrf mode

`VTEP2(config)#mac vrf VRF6`	Configure mac vrf as VRF6
`VTEP2(config-vrf)#rd 2.2.2.2:25`	Configure `route distinguisher value.`
`VTEP2(config-vrf)#route-target both 30.30.30.30:100`	Configure `route-target import and export`
`VTEP2(config-vrf)#exit`	Exit from vrf mode
`VTEP2(config)#mac vrf VRF7`	Configure mac vrf as VRF7
`VTEP2(config-vrf)#rd 2.2.2.2:26`	Configure `route distinguisher value.`
`VTEP2(config-vrf)#route-target both 40.40.40.40:100`	Configure `route-target import and export`
`VTEP2(config-vrf)#exit`	Exit from vrf mode
`VTEP2(config)#mac vrf VRF8`	Configure mac vrf as VRF8
`VTEP2(config-vrf)#rd 2.2.2.2:27`	Configure `route distinguisher value.`
`VTEP2(config-vrf)#route-target both 50.50.50.50:100`	Configure `route-target import and export`
`VTEP2(config-vrf)#exit`	Exit from vrf mode
`VTEP2(config)#mac vrf VRF9`	Configure mac vrf as VRF9
`VTEP2(config-vrf)#rd 2.2.2.2:28`	Configure `route distinguisher value.`
`VTEP2(config-vrf)#route-target both 60.60.60.60:100`	Configure `route-target import and export`
`VTEP2(config-vrf)#exit`	Exit from vrf mode
`VTEP2(config)#mac vrf VRF10`	Configure mac vrf as VRF10
`VTEP2(config-vrf)#rd 2.2.2.2:29`	Configure `route distinguisher value.`
`VTEP2(config-vrf)#route-target both 70.70.70.70:100`	Configure `route-target import and export`
`VTEP2(config-vrf)#exit`	Exit from vrf mode
`VTEP2(config)#nvo vxlan vtep-ip-global 2.2.2.2`	Enable vxlan Source Vtep Ip address global configuration
`VTEP2(config)#nvo vxlan id 10 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF1`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 20 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.

`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF2`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 21 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF3`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 22 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF3`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 23 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF4`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 24 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF5`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 25 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF6`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 26 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF7`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 27 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF8`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 28 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF9`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.

`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#nvo vxlan id 29 ingress-replication`	Enable VxLAN Network Identifier Head End Replication tenant type.
`VTEP2(config-nvo)#vxlan host-reachability-protocol evpn-bgp VRF10`	Host reachability protocol multiprotocol BGP VRF to carry EVPN routes.
`VTEP2(config-nvo)#exit`	Exit from the VxLAN mode.
`VTEP2(config)#qos enable`	Enable QoS.
`VTEP2(config)#hostname VTEP2`	Configure `system's network name as VTEP2.`
`VTEP2(config)#interface lo`	Enter loopback interface mode.
`VTEP2(config-if)#ip address 2.2.2.2/32 secondary`	Configure the secondary IP address of the loopback interface.
`VTEP2(config-if)#exit`	Exit from interface mode.
`VTEP2(config)#interface xe47`	Enter interface mode.
`VTEP2(config-if)#switchport`	Enter the switchport mode.
`VTEP2(config-if)#load-interval 30`	Configure load interval.
`VTEP2(config-if)#exit`	Exit from interface mode.
`VTEP2(config)#interface xe48`	Enter interface mode.
`VTEP2(config-if)#ip address 30.30.30.1/24`	Configure the IP address of the interface.
`VTEP2(config-if)#exit`	Enter interface mode.
`VTEP2(config)#interface xe57`	Enter interface mode.
`VTEP2(config-if)#switchport`	Enter the switchport mode.
`VTEP2(config-if)#load-interval 30`	Configure load interval.
`VTEP2(config-if)#exit`	Exit from interface mode.
`VTEP2(config)#router ospf 100`	Configure router ospf process ID.
`VTEP2(config-router)#ospf router-id 2.2.2.2`	Configure OSPF router ID.
`VTEP2(config-router)#bfd all-interfaces`	Enable BFD all interfaces.
`VTEP2(config-router)#network 2.2.2.2/32 area 0.0.0.0`	Configure network and area as 0.
`VTEP2(config-router)#network 30.30.30.0/24 area 0.0.0.0`	Configure network and area as 0.

VTEP2(config-router)#exit	Exit from router OSPF mode.
VTEP2(config)#router bgp 500	Configure router BGP AS number.
VTEP2(config-router)#bgp router-id 2.2.2.2	Configure BGP router ID.
VTEP2(config-router)#neighbor 1.1.1.1 remote-as 500	Configure a neighbor router and Peer AS Specify AS number of BGP neighbor.
VTEP2(config-router)#neighbor 1.1.1.1 update-source lo	Configure a neighbor router and Source of routing updates as loopback.
VTEP2(config-router)#neighbor 1.1.1.1 advertisement-interval 0	Configure a neighbor router and minimum interval between sending BGP routing updates.
VTEP2(config-router)#address-family ipv4 unicast	Enter Address Family command mode.
VTEP2(config-router-af)#network 2.2.2.2/32	Configure a network to announce via BGP.
VTEP2(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor.
VTEP2(config-router-af)#exit-address-family	Exit from address family mode.
VTEP2(config-router)#address-family l2vpn evpn	Enter Address Family with l2vpn evpn Identifier.
VTEP2(config-router-af)#neighbor 1.1.1.1 activate	Activate the neighbor.
VTEP2(config-router-af)#exit-address-family	Exit from address family mode.
VTEP2(config-router)#exit	Exit from router bgp mode.
VTEP2(config)#nvo vxlan max-cache-disable 2500	Configure vxlan Max number of ARP/ND cache disable allowed for port-vlan.
VTEP2(config)#nvo vxlan access-if port-vlan xe47 2	Configure VxLAN access-if single tagged interface name with VLAN id.
VTEP2(config-nvo-acc-if)#map vnid 22	Map access port attribute with VxLAN Identifier.
VTEP2(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP2(config)#nvo vxlan access-if port-vlan xe47 3	Configure VxLAN access-if single tagged interface name with VLAN id.
VTEP2(config-nvo-acc-if)#map vnid 23	Map access port attribute with VxLAN Identifier.
VTEP2(config-nvo-acc-if)#exit	Exit from access-if mode.
VTEP2(config)#nvo vxlan access-if port-vlan xe47 4	Configure VxLAN access-if single tagged interface name with VLAN id.
VTEP2(config-nvo-acc-if)#map vnid 24	Map access port attribute with VxLAN Identifier.
VTEP2(config-nvo-acc-if)#exit	Exit from access-if mode.

`VTEP2(config)#nvo vxlan access-if port-vlan xe47 5`	Configure VxLAN access-if `single tagged interface name with VLAN id.`
`VTEP2(config-nvo-acc-if)#map vnid 25`	Map access port attribute with VxLAN Identifier.
`VTEP2(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP2(config)#nvo vxlan access-if port-vlan xe47 6`	Configure VxLAN access-if `single tagged interface name with VLAN id.`
`VTEP2(config-nvo-acc-if)#map vnid 26`	Map access port attribute with VxLAN Identifier.
`VTEP2(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP2(config)#nvo vxlan access-if port-vlan xe47 7`	Configure VxLAN access-if `single tagged interface name with VLAN id.`
`VTEP2(config-nvo-acc-if)#map vnid 27`	Map access port attribute with VxLAN Identifier.
`VTEP2(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP2(config)#nvo vxlan access-if port-vlan xe47 8`	Configure VxLAN access-if `single tagged interface name with VLAN id.`
`VTEP2(config-nvo-acc-if)#map vnid 28`	Map access port attribute with VxLAN Identifier.
`VTEP2(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP2(config)#nvo vxlan access-if port-vlan xe47 9`	Configure VxLAN access-if `single tagged interface name with VLAN id.`
`VTEP2(config-nvo-acc-if)#map vnid 29`	Map access port attribute with VxLAN Identifier.
`VTEP2(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP2(config)#nvo vxlan access-if port-vlan xe47 10`	Configure VxLAN access-if `single tagged interface name with VLAN id.`
`VTEP2(config-nvo-acc-if)#map vnid 10`	Map access port attribute with VxLAN Identifier.
`VTEP2(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP2(config)#nvo vxlan access-if port-vlan xe47 11`	Configure VxLAN access-if `single tagged interface name with VLAN id`
`VTEP2(config-nvo-acc-if)#map vnid 21`	Map access port attribute with VxLAN Identifier.
`VTEP2(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP2(config)#nvo vxlan access-if port-vlan xe47 12`	Configure VxLAN access-if `single tagged interface name with VLAN id`
`VTEP2(config-nvo-acc-if)#map vnid 20`	Map access port attribute with VxLAN Identifier.

`VTEP2(config-nvo-acc-if)#exit`	Exit from access-if mode.
`VTEP2(config)#commit`	Commit the candidate configuration to the running configuration.

Validation

Verify OSPF neighbors

```
VTEP1#show ip ospf  neighbor

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID     Pri   State           Dead Time   Address         Interface
Instance ID
11.11.11.11       1   Full/DR         00:00:29    10.10.10.2      xe48                    0
VTEP1#
```

Checking the IP Routes

```
VTEP1#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, 01:21:26
O            2.2.2.2/32 [110/3] via 10.10.10.2, xe48, 01:15:25
C            10.10.10.0/24 is directly connected, xe48, 01:16:11
O            11.11.11.11/32 [110/2] via 10.10.10.2, xe48, 01:15:25
C            20.20.20.0/24 is directly connected, xe52, 01:20:42
O            30.30.30.0/24 [110/2] via 10.10.10.2, xe48, 01:15:25
C            127.0.0.0/8 is directly connected, lo, 01:21:26

Gateway of last resort is not set
VTEP1#
VTEP1#
VTEP1#
```

Verify the BGP neighbors

```
VTEP1#show ip bgp neighbors
BGP neighbor is 2.2.2.2, remote AS 500, local AS 500, internal link, peer index: 12
  BGP version 4, local router ID 1.1.1.1, remote router ID 2.2.2.2
  BGP state = Established, up for 01:15:26
  Last read 00:00:18, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family L2VPN EVPN: advertised and received
  Received 527 messages, 0 notifications, 0 in queue
  Sent 502 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 0 seconds
  Update source is lo

For address family: IPv4 Unicast  BGP table version 2, neighbor version 2
  Index 1, Offset 0, Mask 0x2
  AIGP is enabled
  Community attribute sent to this neighbor (both)
  Large Community attribute sent to this neighbor
  1 accepted prefixes
  1 announced prefixes


For address family: L2VPN EVPN  BGP table version 96, neighbor version 95
  Index 1, Offset 0, Mask 0x2
  Community attribute sent to this neighbor (both)

.skipping 1 line
  31 accepted prefixes
  Accepted AD:0 MACIP:20 MCAST:11 ESI:0 PREFIX:0
  21 announced prefixes

Connections established 1; dropped 0
Local host: 1.1.1.1, Local port: 179
Foreign host: 2.2.2.2, Foreign port: 38227
TCP MSS: (0), Advertise TCP MSS: (1460), Send TCP MSS: (1460),  Receive TCP MSS: (1460)
Sock FD : (22)
Nexthop: 1.1.1.1
Nexthop global: ::
Nexthop local: ::
BGP connection: non shared network
```

Verify the VxLAN access-if

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

                Inner                          Admin   Link
Interface  Vlan  vlan  Ifindex    Vnid        status  status
-----------------------------------------------------------
xe49       2     ---   0x7a120    22          up      up
xe49       3     ---   0x7a121    23          up      up
xe49       4     ---   0x7a122    24          up      up
xe49       5     ---   0x7a123    25          up      up
xe49       6     ---   0x7a124    26          up      up
xe49       7     ---   0x7a125    27          up      up
xe49       8     ---   0x7a126    28          up      up
xe49       9     ---   0x7a127    29          up      up
xe49       10    ---   0x7a128    10          up      up
xe49       11    ---   0x7a129    21          up      up
xe49       12    ---   0x7a12a    20          up      up

Total number of entries are 11
```

Note: Refer sub-interface config for VLAN information.

Verify the VxLAN tunnel

```
VTEP1#
VTEP1#
VTEP1#show nvo vxlan tunnel
VXLAN Network tunnel Entries
Source           Destination      Status        Up/Down       Update
================================================================================
1.1.1.1          2.2.2.2          Installed     01:15:37      01:15:37

Total number of entries are 1
VTEP1#
```

Verify the VxLAN

```
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           2.2.2.2
10       ----          --       AC   xe49      --- Single Homed Port ---     10        -
---      ----              ----
20       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
20       ----          --       AC   xe49      --- Single Homed Port ---     12        -
---      ----              ----
21       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
21       ----          --       AC   xe49      --- Single Homed Port ---     11        -
---      ----              ----
22       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
22       ----          --       AC   xe49      --- Single Homed Port ---     2         -
---      ----              ----
23       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
23       ----          --       AC   xe49      --- Single Homed Port ---     3         -
---      ----              ----
24       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
24       ----          --       AC   xe49      --- Single Homed Port ---     4         -
---      ----              ----
25       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
25       ----          --       AC   xe49      --- Single Homed Port ---     5         -
---      ----              ----
26       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
26       ----          --       AC   xe49      --- Single Homed Port ---     6         -
---      ----              ----
27       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
27       ----          --       AC   xe49      --- Single Homed Port ---     7         -
---      ----              ----
28       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
28       ----          --       AC   xe49      --- Single Homed Port ---     8         -
---      ----              ----
29       ----          L2       NW   ----      ----                          ----      -
---      1.1.1.1           2.2.2.2
29       ----          --       AC   xe49      --- Single Homed Port ---     9         -
---      ----              ----

Total number of entries are 22
```

Note: Refer sub-interface config for VLAN information.

Verify the interface counters

```
VTEP1#

VTEP1#show interface counters rate mbps
+------------------+-------------+------------+-------------+------------+
|     Interface    |    Rx mbps  |    Rx pps  |    Tx mbps  |    Tx pps  |
+------------------+-------------+------------+-------------+------------+
 xe48                42.73         30012        14.25         10011
 xe49                41.60         40625        10.24         10000
 xe54                0.00          0            20.80         20312
VTEP1#
```

Validation for Port Mirroring

Verify the monitor

```
VTEP1#show monitor
Session  State        Reason                 Description
-------  -----------  ---------------------  --------------------------------
1        up           The session is up
VTEP1#

Verify the monitor session

VTEP1#show monitor session 1
   session 1
---------------
type              : local
state             : up
source intf       :
    tx            :
    rx            : xe49
    both          :
source VLANs      :
    rx            :
destination ports : xe54
filter count      : 1

Legend: f = forwarding enabled, l = learning enabled

VTEP1#

VTEP1#show monitor session 1 brief
   session 1
---------------
type              : local
state             : up
source intf       :
    tx            :
```

```
    rx             : xe49
    both           :
destination ports  : xe54
filter count       : 1

VTEP1#

VTEP1#show monitor session 1 filter
   session 1
---------------
filter count       : 1

  ---------------
  match set 1
  ---------------
Sequence number : 10 vlan : 2-6

VTEP1#

END
```

Revised CLI Commands

hardware-profile filter (XGS)

The existing hardware-profile filter CLI syntax is updated as follows:

`hardware-profile filter port-isolation (ingress-ipv4|ingress-ipv6|egress-ipv6|ingress-arp|bfd-group) (enable|disable)`

to

`hardware-profile filter port-isolation (`**`ingress-mirror`**`|ingress-ipv4|ingress-ipv6|egress-ipv6|ingress-arp|bfd-group) (enable|disable)`

Refer to hardware-profile filter (XGS) CLI section for more details.

Use the new filter ingress-mirror profile for port mirroring when monitor session is installed with filters. when the specified filter profile is not enabled, port mirror uses default L2 group.

Command Syntax

```
hardware-profile filter ingress-mirror (disable | enable)
```

Parameter Description

`enable`	Enable the ingress TCAM group for port-mirroring
`disable`	Disable the ingress TCAM group for port-mirroring

Default Value

N/A

Applicability

This command was introduced in OcNOS Version 6.4.1.

Command Mode

Configure mode

Example

```
OcNOS#configure terminal
OcNOS(config)#hardware-profile filter ingress-mirror enable
```

filter

The existing filter CLI syntax is updated as follows:

filter {vlan <2-4094> | cos <0-7> ...

```
(<1-268435453>/<1-4294967294> |) filter {vlan <2-4094>| cos <0-7> | dest-mac (host
XXXX.XXXX.XXXX | XXXX.XXXX.XXXX XXXX.XXXX.XXXX) | src-mac (host XXXX.XXXX.XXXX |
XXXX.XXXX.XXXX XXXX.XXXX.XXXX) | frame-type (ETHTYPE | arp (req | resp|) (sender-ip
A.B.C.D|) (target-ip A.B.C.D|) | ipv4 (src-ip (A.B.C.D | A.B.C.D/M)|) (dest-ip (A.B.C.D
| A.B.C.D/M)|) | ipv6 (src-ip X:X::X:X/M |) (dest-ip X:X::X:X/M |))}
```

to

```
(<1-268435453>/<1-4294967294> |) filter {vlan VLAN_RANGE|inner-vlan VLAN_RANGE| cos <0-
7> | dest-mac (host XXXX.XXXX.XXXX | XXXX.XXXX.XXXX XXXX.XXXX.XXXX) | src-mac (host
XXXX.XXXX.XXXX | XXXX.XXXX.XXXX XXXX.XXXX.XXXX) | frame-type (ETHTYPE | arp (req |
resp|) (sender-ip A.B.C.D|) (target-ip A.B.C.D|) | ipv4 (src-ip (A.B.C.D | A.B.C.D/M)|)
(dest-ip (A.B.C.D | A.B.C.D/M)|) | ipv6 (src-ip X:X::X:X/M |) (dest-ip X:X::X:X/M |))}
```

Refer to filter CLI section for more details.

Command Syntax

```
(<1-268435453>/<1-4294967294> |) filter {vlan VLAN_RANGE| cos <0-7> | dest-mac
  (host XXXX.XXXX.XXXX | XXXX.XXXX.XXXX XXXX.XXXX.XXXX) | src-mac (host
  XXXX.XXXX.XXXX | XXXX.XXXX.XXXX XXXX.XXXX.XXXX) | frame-type (ETHTYPE | arp (req
  | resp|) (sender-ip A.B.C.D|) (target-ip A.B.C.D|) | ipv4 (src-ip (A.B.C.D |
  A.B.C.D/M)|) (dest-ip (A.B.C.D | A.B.C.D/M)|) | ipv6 (src-ip X:X::X:X/M |) (dest-
  ip X:X::X:X/M |))}
no (<1-268435453>/<1-4294967294>) filter
```

Parameters

(<1-268435453>/<1-4294967294> |)	
	Sequence identifier for each rule.
Inner-vLAN	Specify Inner VLAN ID or range(s)
VLAN_RANGE	VLAN ID 2-4094 or range(s): 2-5,10 or 2-5,7-19
<0-7>	COS number
XXXX.XXXX.XXXX	MAC address
ETHTYPE	Ethertype
arp	ARP frames

`req`	Request frames
`resp`	Response frames
`A.B.C.D`	Single IP address
`A.B.C.D/M`	IP addresses with mask
`X:X::X:X/M`	IPv6 addresses with mask

Default

No default value is specified.

Command Mode

Monitor configure mode

Applicability

This command was introduced before OcNOS version 1.3. The VLAN_RANGE option is available from OcNOS version 6.4.0.

Example

```
#configure terminal
(config)#monitor session 3
(config-monitor)#35 filter vlan 10-20,50
```

Abbreviations

Acronym	Expantion
ACL	Access Control List
MAC	Media Access Control
SPAN	Switch Port Analyzer
VLAN	Virtual LAN
VxLAN	Virtual eXtensible Local Area Network

Route Monitor

Overview

Object Tracking provides a mechanism for tracking the reachability status of objects, such as IP status, using Internet Protocol Service Level Agreement (IP SLA). This feature empowers users to monitor the state of these objects and make decisions based on their status. It permits the configuration of multiple track objects on interfaces, delivering flexibility in managing network link status.

Feature Characteristics

Object Tracking establishes a distinct separation between the tracked objects and the actions initiated by a client when there's a change in the state of a tracked object. Users can configure object tracking types as `any` or `all` on the interface, alongside track IDs that specify which statuses to monitor. Modify the interface's link status to either `up` or `down` based on the selected track type and the statuses of the associated track IDs.

When using `Track type all`, the feature performs a `Boolean AND` operation, requiring every object configured on the interface to be in an `up` state for the interface itself to be considered `up`. If any of these objects are not in an `up` state, the interface is set to d`own`.

Conversely, `Track type any` operates as a `Boolean OR` function, necessitating that at least one object configured on the interface must be in an `up` state for the interface to remain `up`. If none of the tracked objects are in an `up` state, the interface is marked as `down`.

Benefits

Users can ensure network reliability by defining specific tracking criteria and actions, allowing them to take appropriate measures when tracked objects experience status change. This contributes to improved network management and performance.

Prerequisites

Before configuring and utilizing Object Tracking, ensure the following prerequisites:

Track IDs: Users must define and configure the track IDs and corresponding objects they want to track for reachability. These track IDs are essential for the feature to work effectively. Deleting all track IDs from the interface will bring the interface up if it was previously down.

Interface Configuration: The feature involves configuring track types on interfaces. Therefore, ensuring that the interfaces are correctly configured and operational is important. In cases where an interface has both object tracking configurations and next-hop reachability, deleting the object tracking configurations is necessary to bring the interface back up if it goes down.

Object Tracking Criteria: Define the specific criteria and conditions for tracking an object’s reachability, such as IP status, using IP SLA.

Configuration

The below topology illustrates a network configuration involving three routers, R1, R2, and R3, with a central device referred to as the Device Under Test (DUT) positioned in the middle. This topology represents a linear or sequential network structure that showcases the Route Monitor feature.

Topology

A series of configurations were implemented on routers R1, R2, and R3, as well as on the DUT, to showcase the functionality of the Route Monitor feature. The objective was to demonstrate the configuration of network routers to monitor the reachability status of specific IPv4 and IPv6 addresses using IP SLA and illustrate that these configurations can work in conjunction with the Route Monitor feature to enable informed decisions based on the reachability status of tracked objects.

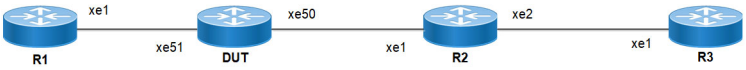

Route Monitor Topology

IPv4 Configuration

DUT

Use the following configuration to set up an IP SLA and enable object tracking on a network device. These commands assign IPv4 addresses to interfaces, configure specific IP SLA parameters such as threshold, timeout, and frequency, create a time-range for scheduling measurements, and establish static routes with nexthop addresses. Configure object tracking to monitor the reachability of tracked objects. These configurations highlight the versatility and functionality of the network device by allowing it to monitor IPv4 addresses, make decisions based on object tracking, and optimize network operations.

`DUT#configure terminal`	Enter configure mode.
`DUT(config)#interface xe50`	Enter interface mode `xe50`.
`DUT(config-if)#ip address 2.2.2.1/24`	Assign the IP address `2.2.2.1` with a subnet mask of `/24` to interface `xe50`.
`DUT(config-if)#exit`	Exit interface mode `xe50`.
`DUT(config)#interface xe51`	Enter interface mode `xe51`.
`DUT(config-if)#ip address 1.1.1.2/24`	Assign the IP address 1.1.1.2 with a subnet mask of /24 to interface xe51.
`DUT(config-if)#exit`	Exit interface mode `xe51`.
`DUT(config)#ip sla 1`	Create an IP SLA operation with index `1`.
`DUT(config-ip-sla)#icmp-echo ipv4 3.3.3.1 source-interface xe50`	Configure the SLA to send ICMP echo requests to destination IPv4 address `3.3.3.1` using interface `xe50` as the source.
`DUT(config-ip-sla-echo)#threshold 1000`	Set the threshold value for SLA to `1000` milliseconds.
`DUT(config-ip-sla-echo)#timeout 1000`	Set the timeout value for SLA to `1000` milliseconds.
`DUT(config-ip-sla-echo)#frequency 5`	Configure the frequency value for SLA to send ICMP echo packets every 5 seconds.
`DUT(config-ip-sla-echo)#exit`	Exit IP SLA echo mode.

`DUT(config-ip-sla)#exit`	Exit IP SLA mode.
`DUT(config)#time-range tr1`	Create a time range named `tr1`.
`DUT(config-tr)#start-time 11:22 3 july 2021`	Set the start time for the time range to `11:22 on July 3, 2021`.
`DUT(config-tr)#end-time after 200`	Set the end time to be `200` minutes from the start time.
`DUT(config-tr)#exit`	Exit time-range mode.
`DUT(config)#ip sla schedule 1 time-range tr1`	Schedule IP SLA operation 1 to run within the specified time range `tr1`.
`DUT(config)#track 1 ip sla 1 reachability`	Creating a tracking object to monitor the reachability status of IP SLA operation `1`.
`DUT(config-object-track)#exit`	Exit object track mode.
`DUT(config)#ip route 3.3.3.0/24 2.2.2.2 track 1`	Add a static route for the destination network `3.3.3.0/24` with next-hop IP `2.2.2.2`, tracked by tracking object `1`.
`DUT(config)#ip route 5.5.5.0/24 1.1.1.2`	Add a static route for the destination network `5.5.5.0/24` with next-hop IP `1.1.1.2`.
`DUT(config)#ip route 6.6.6.0/24 2.2.2.2 track 1`	Add a static route for the destination network `6.6.6.0/24` with next-hop IP `2.2.2.2`, tracked by tracking object `1`.
`DUT(config)#ip route 6.6.6.0/24 1.1.1.2 10`	Add a static route for the destination network `6.6.6.0/24` with next-hop IP `1.1.1.2` and a metric of `10`.
`DUT(config)#commit`	Commit the candidate configuration to the running configuration.
`DUT(config)#interface xe51`	Enter interface mode `xe51`.
`DUT(config-if)#object-tracking all`	Enable object tracking for `all` tracking objects on interface `xe51`.
`DUT(config-if)#object-tracking 1`	Configure object tracking `1` on interface `xe51`.
`DUT(config-if)#object-tracking 2`	Configure object tracking `2` on interface `xe51`.
`DUT(config-if)#exit`	Exit interface mode.
`DUT(config)#exit`	Exit configure mode.

By configuring the routes below, R1, R2, and R3 effectively forward network traffic to its designated destinations within the network. These configurations actively contribute to efficient routing operations and ensure network traffic reaches its targets.

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface xe1`	Enter interface mode `xe1`.
`R1(config-if)#ip address 1.1.1.1/24`	Assign the IP address `1.1.1.1` with a subnet mask of `/24` to interface `xe1`.
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exit interface mode `xe1`.
`R1(config)#ip route 2.2.2.0/24 1.1.1.2`	Add a static route for the destination network `2.2.2.0/24` with next-hop IP `1.1.1.2`.
`R1(config)#ip route 3.3.3.0/24 1.1.1.2`	Add a static route for the destination network `3.3.3.0/24` with next-hop IP `1.1.1.2`.

`R1(config)#commit`	Commit the candidate configuration to the running configuration.
`R1(config)#exit`	Exit configure mode.

R2

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface xe1`	Enter interface mode `xe1`.
`R2(config-if)#ip address 2.2.2.2/24`	Assign the IP address `2.2.2.2` with a subnet mask of `/24` to interface `xe1`.
`R2(config-if)#exit`	Exit interface mode `xe1`.
`R2(config)#interface xe2`	Enter interface mode `xe2`.
`R2(config-if)#ip address 3.3.3.1/24`	Assign the IP address `3.3.3.1` with a subnet mask of `/24` to interface `xe2`.
`R2(config-if)#exit`	Exit interface mode `xe2`.
`R2(config)#ip route 1.1.1.0/24 2.2.2.1`	Add a static route for the destination network `1.1.1.0/24` with next-hop IP `2.2.2.1`.
`R2(config)#commit`	Commit the candidate configuration to the running configuration.
`R2(config)#exit`	Exit configure mode.

R3

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface xe1`	Enter interface mode `xe1`.
`R3(config-if)#ip address 3.3.3.2/24`	Assign the IP address `3.3.3.2` with a subnet mask of `/24` to interface `xe1`.
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exit interface mode `xe1`.
`R3(config)#ip route 1.1.1.0/24 3.3.3.1`	Add a static route for the destination network `1.1.1.0/24` with next-hop IP `3.3.3.1`.
`R3(config)#ip route 2.2.2.0/24 3.3.3.1`	Add a static route for the destination network `2.2.2.0/24` with next-hop IP `3.3.3.1`.
`R3(config)#commit`	Commit the candidate configuration to the running configuration.
`R3(config)#exit`	Exit configure mode.

Validation

The following show output displays information about the IPv4 route table, IP SLA reachability tracking, and interface status on a network device running OcNOS.

DUT

```
DUT#show track
TRACK Id: 1
  IP SLA 1 reachability
```

```
  Reachability is UP
    4 changes, last change : 2019 Mar 14 14:53:47
Track interface : xe51

DUT#show ip route track-table
 ip route 3.3.3.0 255.255.255.0 2.2.2.2 track 1 state is [up]
 ip route 6.6.6.0 255.255.255.0 2.2.2.2 track 1 state is [up]

DUT#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            1.1.1.0/24 is directly connected, xe51, 00:55:38
C            2.2.2.0/24 is directly connected, xe50, 00:49:50
S            3.3.3.0/24 [1/0] via 2.2.2.2, xe50, 00:00:03
S            5.5.5.0/24 [1/0] via 1.1.1.2, xe51, 00:08:12
S            6.6.6.0/24 [1/0] via 2.2.2.2, xe50, 00:00:03

Gateway of last resort is not set

DUT#show interface brief xe51

Codes: ETH - Ethernet, LB - Loopback, AGG - Aggregate, MLAG - MLAG Aggregate
       FR - Frame Relay, TUN -Tunnel, PBB - PBB Logical Port, VP - Virtual Port
       CVP - Channelised Virtual Port, METH - Management Ethernet, UNK- Unknown
       ED - ErrDisabled, PD - Protocol Down, AD - Admin Down, IA - InActive
       PD(Min L/B) - Protocol Down Min-Links/Bandwidth
       OTD  - Object Tracking Down
       DV - DDM Violation, NA - Not Applicable
       NOM - No operational members, PVID - Port Vlan-id
       Ctl - Control Port (Br-Breakout/Bu-Bundle)


--------------------------------------------------------------------------------
Ethernet  Type    PVID   Mode     Status  Reason  Speed Port ch#  Ctl Br/Bu  Loopbk Interface
--------------------------------------------------------------------------------
xe51       ETH    --    routed   down    OTD     10g   --      No       No
```

IPv6 Configuration

DUT

Use the following configuration to set up an IP SLA and enable object tracking on a network device. These commands assign IPv6 addresses to interfaces, configure specific IP SLA parameters such as threshold, timeout, and frequency, create a time-range for scheduling measurements, and establish static routes with nexthop addresses. Configure object tracking to monitor the reachability of tracked objects. These configurations highlight the versatility and functionality of the network device by allowing it to monitor IPv6 addresses, make decisions based on object tracking, and optimize network operations.

`DUT#configure terminal`	Enter configure mode.
`DUT(config)#interface xe50`	Enter interface mode `xe50`.
`DUT(config-if)#ipv6 address 2000::1/64`	Assign an IPv6 address (`2000::1/64`) to interface `xe50`.
`DUT(config-if)#exit`	Exit interface mode `xe50`.
`DUT(config)#interface xe51`	Enter interface mode `xe51`.
`DUT(config-if)#ipv6 address 1000::2/64`	Assign an IPv6 address (`1000::2/64`) to interface `xe51`.

`DUT(config-if)#exit`	Exit interface mode `xe51`.
`DUT(config)#ip sla 1`	Create an IP SLA operation with index `1`.
`DUT(config-ip-sla)#icmp-echo ipv6 3000::1 source-interface xe50`	Configure the SLA to send IPv6 ICMP echo requests to destination IPv6 address `3000::1` using interface `xe50` as the source.
`DUT(config-ip-sla-echo)#threshold 1000`	Set the threshold value for SLA to `1000` milliseconds.
`DUT(config-ip-sla-echo)#timeout 1000`	Set the timeout value for SLA to `1000` milliseconds.
`DUT(config-ip-sla-echo)#frequency 5`	Configure the frequency value for SLA to send IPv6 ICMP echo packets every 5 seconds.
`DUT(config-ip-sla-echo)#exit`	Exit IP SLA echo mode.
`DUT(config-ip-sla)#exit`	Exit IP SLA mode.
`DUT(config)#time-range tr1`	Create a time range named `tr1`.
`DUT(config-tr)#start-time 11:22 3 july 2021`	Set the start time for the time range to `11:22 on July 3, 2021`.
`DUT(config-tr)#end-time after 200`	Set the end time to be `200` minutes from the start time.
`DUT(config-tr)#exit`	Exit time-range mode.
`DUT(config)#ip sla schedule 1 time-range tr1`	Schedule IP SLA operation `1` to run within the specified time range `tr1`.
`DUT(config)#track 1 ip sla 1 reachability`	Creating a tracking object to monitor the reachability status of IP SLA operation `1`.
`DUT(config-object-track)#exit`	Exit object track mode.
`DUT(config)#ipv6 route 3000::0/64 2000::2 track 1`	Add an IPv6 static route for the destination network `3000::0/64` with a next-hop IPv6 `2000::2`, tracked by tracking object `1`.
`DUT(config)#ipv6 route 3333::1/128 1000::1`	Add an IPv6 static route for the destination network `3333::1/128` with next-hop IPv6 `1000::1`.
`DUT(config)#ipv6 route 3333::1/128 2000::2 track 1`	Add an IPv6 static route for the destination network `6.6.6.0/24` with next-hop IPv6 `2000::2`, tracked by tracking object `1`.
`DUT(config)#ipv6 route 3333::1/128 1000::1 10`	Add an IPv6 static route for the destination network `3333::1/128` with next-hop IP `1000::1` and a metric of `10`.
`DUT(config)#commit`	Commit the candidate configuration to the running configuration.
`DUT(config)#interface xe51`	Enter interface mode `xe51`.
`DUT(config-if)#object-tracking all`	Enable object tracking for `all` tracking objects on interface `xe51`.
`DUT(config-if)#object-tracking 1`	Configure object tracking `1` on interface `xe51`.
`DUT(config-if)#object-tracking 2`	Configure object tracking `2` on interface `xe51`.
`DUT(config-if)#exit`	Exit interface mode.
`DUT(config)#exit`	Exit configure mode.

By configuring the routes below, R1, R2, and R3 effectively forward network traffic to its designated destinations within the network. These configurations actively contribute to efficient routing operations and ensure network traffic reaches its targets.

R1

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface xe1`	Enter interface mode `xe1`.
`R1(config-if)#ipv6 address 1000::1/64`	Assign the IPv6 address `1000::1` with a subnet mask of `/64` to interface `xe1`.
`R1(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R1(config-if)#exit`	Exit interface mode `xe1`.
`R1(config)#ipv6 route 2000::0/64 1000::2`	Add an IPv6 static route for the destination network `2000::0/64` with next-hop IPv6 `1000::2`.
`R1(config)#ipv6 route 3000::0/64 1000::2`	Add an IPv6 static route for the destination network `3000::0/64` with next-hop IPv6 `1000::2`.
`R1(config)#commit`	Commit the candidate configuration to the running configuration.
`R1(config)#exit`	Exit configure mode.

R2

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface xe1`	Enter interface mode `xe1`.
`R2(config-if)#ipv6 address 2000::2/64`	Assign the IPv6 address `2000::2` with a subnet mask of `/64` to interface `xe1`.
`R2(config-if)#exit`	Exit interface mode `xe1`.
`R2(config)#interface xe2`	Enter interface mode `xe2`.
`R2(config-if)#ipv6 address 3000::1/64`	Assign the IPv6 address `3000::1` with a subnet mask of `/64` to interface `xe2`.
`R2(config-if)#exit`	Exit interface mode `xe2`.
`R2(config)#ipv6 route 1000::0/64 2000::1`	Add an IPv6 static route for the destination network `1000::0/64` with next-hop IPv6 `2000::1`.
`R2(config)#commit`	Commit the candidate configuration to the running configuration.
`R2(config)#exit`	Exit configure mode.

R3

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface xe1`	Enter interface mode `xe1`.
`R3(config-if)#ipv6 address 3000::2/64`	Assign the IPv6 address `3000::2` with a subnet mask of `/64` to interface `xe1`.
`R3(config-if)#commit`	Commit the candidate configuration to the running configuration.
`R3(config-if)#exit`	Exit interface mode `xe1`.
`R3(config)#ipv6 route 1000::0/64 3000::1`	Add an IPv6 static route for the destination network `1000::0/64` with next-hop IPv6 `3000::1`.

`R3(config)#ipv6 route 2000::0/64 3000::1`	Add an IPv6 static route for the destination network `2000::0/64` with next-hop IPv6 `3000::1`.
`R3(config)#commit`	Commit the candidate configuration to the running configuration.
`R3(config)#exit`	Exit configure mode.

Validation

The following show output displays the information about IP SLA reachability tracking, IPv6 route tables, and interface status on a network device running OcNOS.

DUT

```
DUT#show track
TRACK Id: 1
  IP SLA 1 reachability
  Reachability is UP
    4 changes, last change : 2019 Mar 14 14:53:47
Track interface : xe51

DUT#show ip route track-table
 ipv6 route 3000::0/64 2000::2 track 1 state is [up]
 ipv6 route 3333::1/128 2000::2 track 1 state is [up]

DUT#show ip sla summary
IP SLA Operation Summary
Codes: * active, ^ inactive

ID       Type        Destination       Stats       Return      Last
                                       (usec)        Code        Run
---------------------------------------------------------------------
*1      icmp-echo   3000::1                16000         OK        2019 Mar 11 1
6:11:40
---------------------------------------------------------------------

DUT#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, 00:04:46
C       1000::/64 via ::, xe51, 00:02:48
C       2000::/64 via ::, xe50, 00:02:48
S       3000::/64 [1/0] via 2000::2, xe50, 00:02:48
S       3333::1/128 [1/0] via 2000::2, xe50, 00:02:48

DUT#show interface brief xe51

Codes: ETH - Ethernet, LB - Loopback, AGG - Aggregate, MLAG - MLAG Aggregate
       FR - Frame Relay, TUN -Tunnel, PBB - PBB Logical Port, VP - Virtual Port
       CVP - Channelised Virtual Port, METH - Management Ethernet, UNK- Unknown
       ED - ErrDisabled, PD - Protocol Down, AD - Admin Down, IA - InActive
       PD(Min L/B) - Protocol Down Min-Links/Bandwidth
       OTD  - Object Tracking Down
       DV - DDM Violation, NA - Not Applicable
       NOM - No operational members, PVID - Port Vlan-id
       Ctl - Control Port (Br-Breakout/Bu-Bundle)


--------------------------------------------------------------------------------
Ethernet   Type     PVID  Mode      Status  Reason  Speed Port Ch #  Ctl Br/Bu  Loopbk
Interface
```

```
-------------------------------------------------------------------------------------
xe51        ETH     --    routed     down    OTD     10g   --         No         No
```

Implementation Examples

Here is a practical scenario and use case for Object Tracking implementation:

Link Redundancy: Object Tracking can be used to monitor the reachability of primary and backup network links. If the primary link fails or becomes congested, the system can automatically switch traffic to the backup link, ensuring uninterrupted network connectivity.

Load Balancing: Object Tracking helps optimize load balancing by continuously assessing the health and availability of servers or paths. If a server becomes overloaded or fails, traffic can be intelligently redirected to healthy servers, improving resource utilization and user experience.

Failover Testing and Verification: Object Tracking provides a mechanism for simulating network failures and verifying failover mechanisms. By configuring tracked objects to mimic real-world conditions, network administrators can assess the resilience of their network configurations and ensure they perform as expected during failures.

New CLI Commands

The Route Monitor feature introduces the following configuration commands. For more information, refer to the Interface Commands, IP Service Level Agreements Commands, and Object Tracking Commands chapters in the System Management Guide, Release 6.4.1.

object-tracking

Use this command to configure track IDs and options on the interfaces.

Use the no parameter with this command to remove the configurations.

These commands configure object tracking on interfaces, with specific track IDs and tracked objects set to determine what gets tracked and affects the interface's status.

The `object-tracking` command provides flexibility, enabling both `all` and `any` tracking behaviors for influencing the interface's status. A maximum of 8 track IDs can be configured per interface. It is possible to configure the same track IDs or options on multiple interfaces.

Command Syntax

```
object-tracking <1-500>
object-tracking <all | any>
no object-tracking <1-500>
no object-tracking <all | any>
```

Parameters

`<1-500>`	Object tracking ID
`all`	Boolean AND operation. Each object configured on the interface must be in an up state for the interface itself to be in an up state; otherwise, it will be brought down.
`any`	Boolean OR operation. At least one object configured on the interface must be in an up state; otherwise, the interface will be brought down.

Default

None

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Example

Here are some example commands for configuring object tracking in the interface mode.

```
OcNOS(config)#int xe5
OcNOS(config-if)#object-tracking 10
OcNOS(config-if)#object-tracking all
OcNOS(config-if)#commit

OcNOS(config-if)#no object-tracking 10
OcNOS(config-if)#no object-tracking all
OcNOS(config-if)#commit
OcNOS(config-if)#exit
```

Troubleshooting

Interface Status: Verify the status of the interface linked with object tracking. If the configured track type is `all`, confirm that all tracked objects are in an `up` state to consider the interface as `up`. In the case of the track type being `any`, ensure that at least one tracked object is `up` to maintain the interface in an `up` state.

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
NSM	Network and Service Management
IP SLA	Internet Protocol Service Level Agreement
DUT	Device Under Test

Glossary

The following provides definitions for key terms used throughout this document.

Object Tracking	A feature that monitors the reachability status of objects, such as IP status, using IP SLA and allows users to take actions based on their status.
Track Object	An object configured for tracking within the Object Tracking feature. These objects can represent specific network components or conditions, such as IP addresses or link statuses.
Track ID	A unique identifier associated with a track object that enables the system to monitor and assess the status of that object.
Track Type	The configuration specifies how the interface's link status should be determined based on the statuses of associated track objects. It can be set to all or any.
Track Type “All”	A track type that uses a Boolean AND function, requiring that all tracked objects be in an up state for the interface to be considered up.
Track Type “Any”	A track type that uses a Boolean OR function, ensuring that at least one tracked object is in an up state for the interface to remain up.

DHCP Server Group

Overview

Dynamic Host Control Protocol (DHCP) Group provides the capability to specify multiple DHCP servers as a group on the DHCP relay agent and to correlate a relay agent interface with the server group. When the interface receives request messages from clients, the relay agent forwards the message to all the DHCP servers of the group. One or multiple DHCP servers in the group process the request and respond with an offer to the client. The client reviews the offer and sends the request message to the chosen server to obtain the network configuration that includes an IP address. The illustration below shows a DHCP client sending a request message to a DHCP relay agent that forwards the message to the three servers in the DHCP server group to get their network configuration. The DHCP client and DHCP relay agent run OcNOS, but the DHCP servers can be OcNOS or Linux devices.

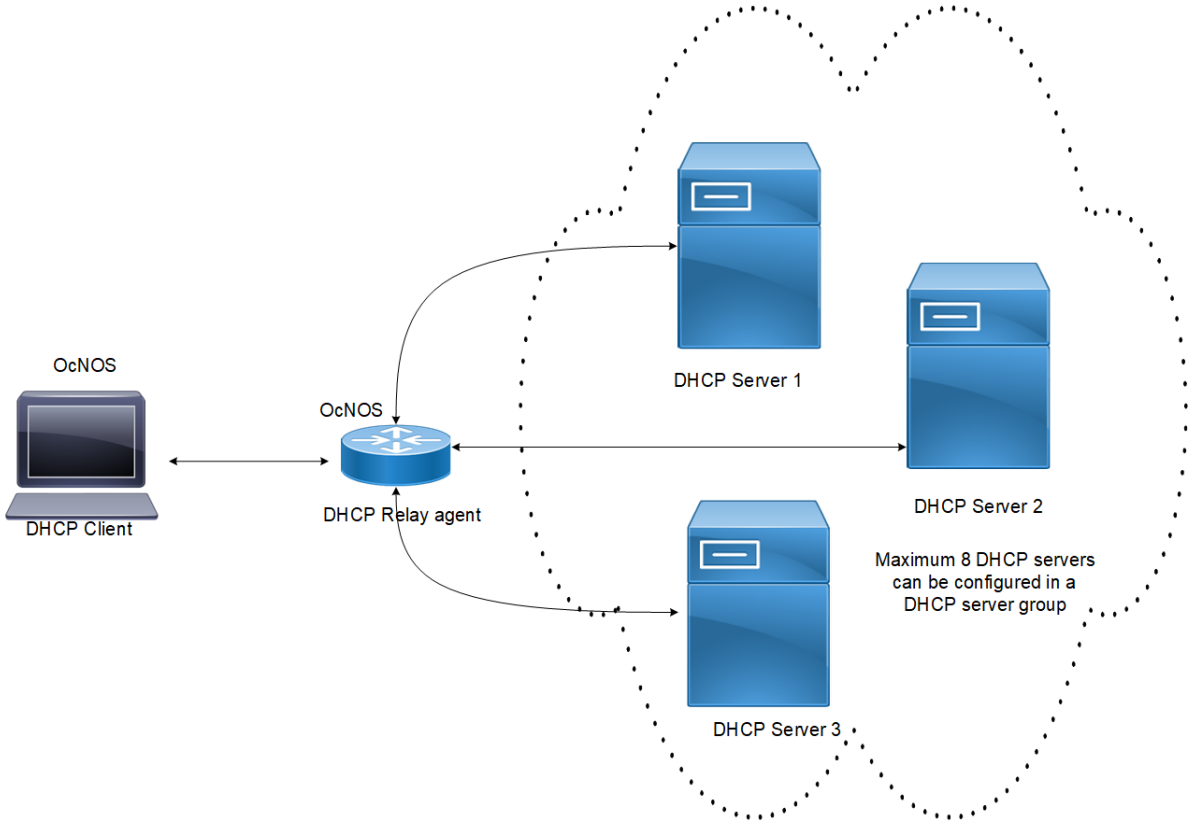

DHCP server group

Feature Characteristics

This feature enables the configuration of the DHCP server group and attaches it to a DHCP relay agent through the CLI and the NetConf interface. A DHCP server group can be attached with multiple DHCP relay uplink interfaces, but at a given time, a single DHCP relay uplink interface is allowed to be attached with a single DHCP server group. The attachment of the DHCP relay uplink interface to another DHCP server group dissociates its attachment with the earlier attached DHCP server group.

This feature helps to configure DHCP IPv4 and IPv6 groups and attach server IP addresses to the group. Creating a maximum of 32 IPv4 and 32 IPv6 groups per VRF is allowed, and configuring 8 DHCP servers is permitted for each DHCP server group.

Benefits

The DHCP relay agent forwards the request message from the DHCP client to multiple DHCP servers in the group. Forwarding the request message to multiple DHCP servers increases the reliability of obtaining the network configuration.

Configuration

Before configuring the DHCP client and the DHCP relay agent, make sure that DHCP server is configured and the `dhcpd` service is running in the DHCP server.

Topology

In the below example, DHCP server1 and DHCP server2 (OcNOS or Linux devices) are connected to the DHCP relay agent (an OcNOS device), and the DHCP relay is connected to a DHCP client (an OcNOS device). The DHCP client sends discover message to the DHCP servers through the DHCP relay agent.

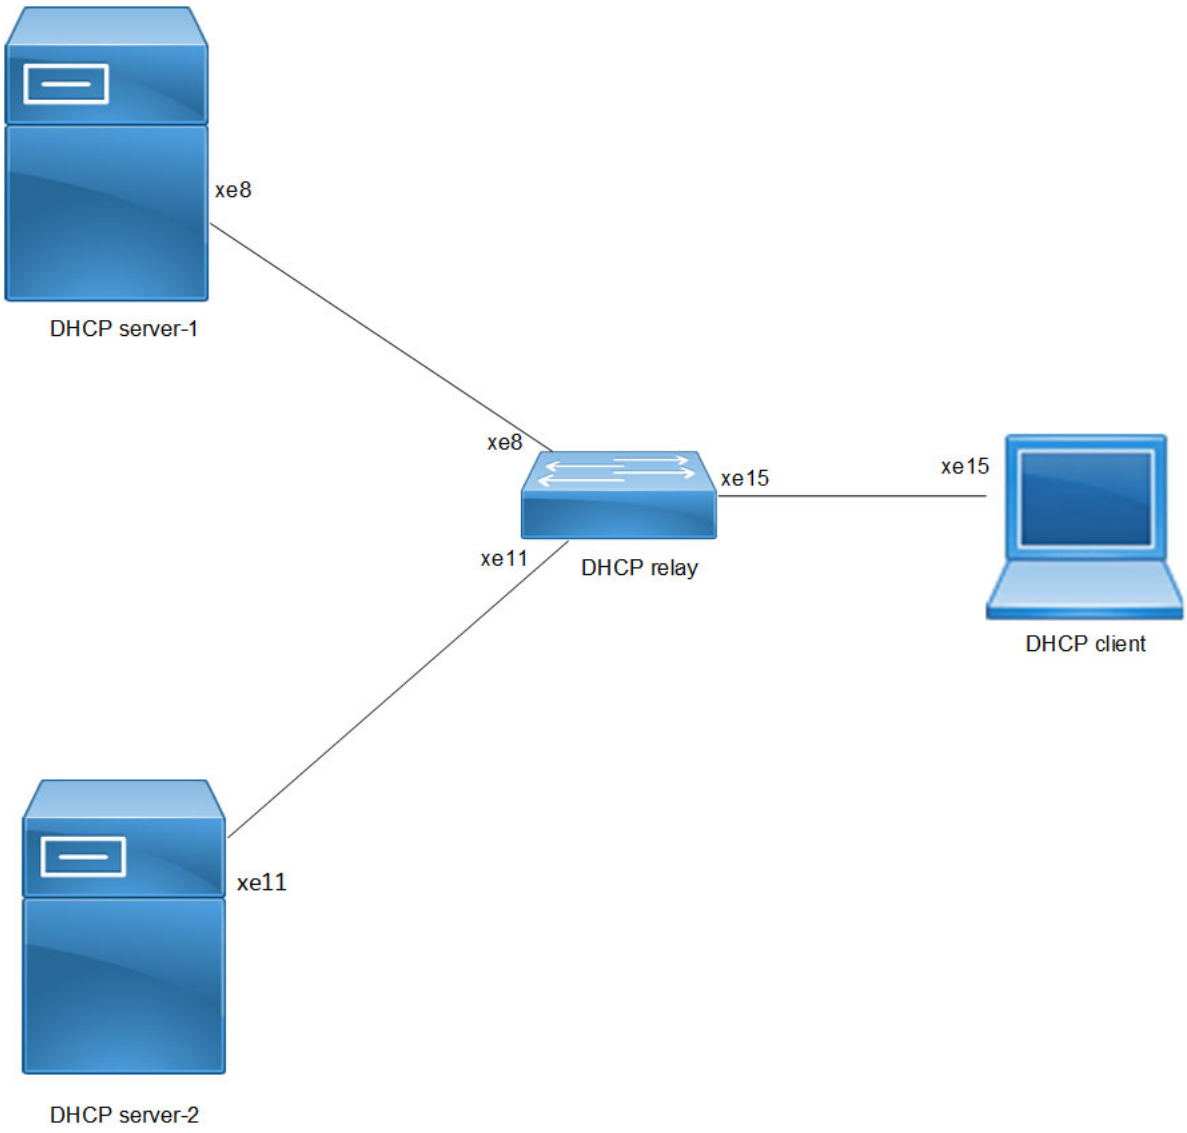

DHCP server group topology

Configuration

DHCP Client Configuration for IPv4

`OcNOS#configure terminal`	Enter configure mode.
`OcNOS(config)#feature dhcp`	Enable the feature DHCP. This will be enabled by default.
`OcNOS(config)#int xe15`	Enter interface mode `xe15`.
`OcNOS(config-if)#ip address dhcp`	The client requests for the IP address to the server, once it receives the acknowledgment from the server, it assigns the IP address to the interface in which this command is enabled.

`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.

Validation

The below shows the running configuration of the DHCPv4 client node:

```
OcNOS#show running-config dhcp
interface eth0
 ip address dhcp
!
interface xe15
 ip address dhcp

OcNOS#show ip interface brief

'*' - address is assigned by dhcp client

Interface            IP-Address      Admin-Status          Link-Status
cd1                  unassigned      up                    down
cd3                  unassigned      up                    down
ce0                  unassigned      up                    down
ce2                  unassigned      up                    down
eth0                *10.12.121.156   up                    up
lo                   127.0.0.1       up                    up
lo.management        127.0.0.1       up                    up
xe4                  unassigned      up                    down
xe5                  unassigned      up                    down
xe6                  unassigned      up                    down
xe7                  unassigned      up                    down
xe8                  unassigned      up                    down
xe9                  unassigned      up                    down
xe10                 unassigned      up                    down
xe11                 unassigned      up                    down
xe12                 unassigned      up                    down
xe13                 unassigned      up                    down
xe14                 unassigned      up                    down
xe15                *20.20.20.1      up                    up
xe16                 unassigned      up                    down
xe17                 unassigned      up                    down
xe18                 unassigned      up                    down
xe19                 unassigned      up                    down
xe20                 unassigned      up                    down
xe21                 unassigned      up                    down
xe22                 unassigned      up                    down
xe23                 unassigned      up                    down
xe24                 unassigned      up                    down
xe25                 unassigned      up                    down
xe26                 unassigned      up                    down
xe27                 unassigned      up                    down
OcNOS#--
OcNOS#
OcNOS#show ip int xe15 br

'*' - address is assigned by dhcp client
```

```
Interface            IP-Address      Admin-Status          Link-Status
xe15                *20.20.20.1      up                    up
OcNOS#
```

DHCP Relay Agent Configuration for IPv4

`OcNOS#configure terminal`	Enter configure mode.
`OcNOS(config)#ip dhcp relay server-group dhcp-relay-gp`	Configure relay server-group group name in global mode.
`OcNOS(dhcp-relay-group)#server 10.10.10.2`	Configure server `10.10.10.2`.
`OcNOS(dhcp-relay-group)#exit`	Exit DHCP relay group.
`OcNOS(config)#interface xe15`	Enter interface mode `xe15`.
`OcNOS(config-if)#ip address 20.20.20.2/24`	Configure IPv4 address `20.20.20.2` on the interface `xe15`.
`OcNOS(config-if)#ip dhcp relay`	Relay should be configured on the interface connecting to the client.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.
`OcNOS(config)#interface xe8`	Enter interface mode `xe8`.
`OcNOS(config-if)#ip address 10.10.10.3/24`	Configure IPv4 address `10.10.10.3` on the interface `xe8`.
`OcNOS(config-if)#ip dhcp relay uplink`	Configure relay uplink on the device connecting the server.
`OcNOS(config-if)#ip dhcp relay server-select dhcp-relay-gp`	Configure relay server-select group name on the device connected to the server.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.
`OcNOS(config)#ip dhcp relay server-group dhcp-relay-gp`	Configure relay server-group group name in global mode.
`OcNOS(dhcp-relay-group)#server 40.10.10.2`	Configure IPv4 DHCP server address `40.10.10.2` on the server group.
`OcNOS(dhcp-relay-group)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(dhcp-relay-group)#exit`	Exit DHCP relay group.
`OcNOS(config)#interface xe11`	Enter interface mode `xe11`.
`OcNOS(config-if)#ip address 40.10.10.3/24`	Configure IPv4 address `40.10.10.3` on the interface `xe11`.
`OcNOS(config-if)#ip dhcp relay uplink`	Configure relay uplink on the device connecting the server.
`OcNOS(config-if)#ip dhcp relay server-select dhcp-relay-gp`	Configure relay server-select group name on the device connected to the server.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.

Validation

The below shows the running configuration of the DHCPv4 relay agent node:

```
OcNOS#show running-config dhcp
interface eth0
 ip address dhcp
!
!

ip dhcp relay server-group dhcp-relay-gp
 server 10.10.10.2
 server 40.10.10.2
interface xe8
 ip dhcp relay uplink
 ip dhcp relay server-select dhcp-relay-gp
!
interface xe11
 ip dhcp relay uplink
 ip dhcp relay server-select dhcp-relay-gp
!
interface xe15
 ip dhcp relay
!
OcNOS#
OcNOS#
OcNOS#show ip dhcp relay
DHCP relay service is Enabled.
VRF Name: default
  Option 82: Disabled
  Interface                Uplink/Downlink
  ---------                -------------
  xe8                      Uplink
  xe11                     Uplink
  xe15                     Downlink
  Interface                Group-Name              Server
  ---------                -------------           ------
  xe11                     dhcp-relay-gp     10.10.10.2,40.10.10.2
Incoming DHCPv4 packets which already contain relay agent option are FORWARDED
u
nchanged.
OcNOS#
```

DHCP Server-1 Configuration for IPv4

`OcNOS#configure terminal`	Enter configure mode.
`OcNOS(config)ip dhcp server pool DHCP-Server-1`	Configure DHCP server group for server in global mode.
`OcNOS(dhcp-config)#network 10.10.10.0 netmask 255.255.255.0`	Configure network `10.10.10.0` and netmask `255.255.255.0`.
`OcNOS(dhcp-config)#address range low-address 10.10.10.1 high-address 10.10.10.254`	Configure address range from `10.10.10.1` to `10.10.10.254`.
`OcNOS(dhcp-config)#dns-server 192.2.2.2`	Configure the DNS server `192.2.2.2`.

`OcNOS(dhcp-config)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(dhcp-config)#exit`	Exit DHCP config mode.
`OcNOS(config)#ip dhcp server pool DHCP-SER`	Configure DHCP server group for client in global mode.
`OcNOS(dhcp-config)#network 20.20.20.0 netmask 255.255.255.0`	Configure network `20.20.20.0` and netmask `255.255.255.0`.
`OcNOS(dhcp-config)#address range low-address 20.20.20.1 high-address 20.20.20.30`	Configure address range from `20.20.20.1` to `20.20.20.30`.
`OcNOS(dhcp-config)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(dhcp-config)#exit`	Exit dhcp config mode.
`OcNOS(config)#interface xe8`	Enter interface mode `xe8`.
`OcNOS(config-if)#ip address 10.10.10.2/24`	Configure IP address on the interface `xe8`.
`OcNOS(config-if)#ip dhcp server`	Server should be configured on the interface while connected to the relay.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.
`OcNOS(config)#ip route 20.20.20.0/24 10.10.10.3`	Configure static route of `20.20.20.0/24` by next hop interface `10.10.10.3`.
`OcNOS(config)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config)#exit`	Exit config mode.

Validation

The below shows the running configuration of the DHCPv4 Server-1 node:

```
OcNOS#show running-config dhcp
interface eth0
 ip address dhcp
!
!

ip dhcp server pool DHCP-Server-1
 network 10.10.10.0 netmask 255.255.255.0
 address range low-address 10.10.10.1 high-address 10.10.10.254
 dns-server 192.2.2.2
ip dhcp server pool DHCP-SER
 network 20.20.20.0 netmask 255.255.255.0
 address range low-address 20.20.20.1 high-address 20.20.20.30
interface xe8
 ip dhcp server
!
OcNOS#
```

DHCP Server-2 Configuration for IPv4

`OcNOS#configure terminal`	Enter configure mode.
`OcNOS(config)#ip dhcp server pool DHCP-Server-2`	Configure DHCP server group for server in global mode.
`OcNOS(dhcp-config)#network 40.10.10.0 netmask 255.255.255.0`	Configure network `40.10.10.0` and netmask `255.255.255.0`.
`OcNOS(dhcp-config)#address range low-address 40.10.10.1 high-address 40.10.10.254`	Configure address range from `40.10.10.1` to `40.10.10.254`.
`OcNOS(dhcp-config)#dns-server 192.2.2.2`	Configure DNS server `192.2.2.2`.
`OcNOS(dhcp-config)#ip dhcp server pool DHCP-SER`	Configure DHCP server group for client in global mode.
`OcNOS(dhcp-config)#network 20.20.20.0 netmask 255.255.255.0`	Configure network `20.20.20.0` and netmask `255.255.255.0`.
`OcNOS(dhcp-config)#address range low-address 20.20.20.1 high-address 20.20.20.30`	Configure address range from `20.20.20.1` to `20.20.20.30`.
`OcNOS(dhcp-config)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(dhcp-config)#exit`	Exit DHCPv6 config mode.
`OcNOS(config)#interface xe11`	Enter interface mode `xe11`.
`OcNOS(config-if)#ip address 40.10.10.2/24`	Configure IP address `40.10.10.2/24` on the interface `xe11`.
`OcNOS(config-if)#ip dhcp server`	Server should be configured on the interface while connected to the relay.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.
`OcNOS(config)#ip route 20.20.20.0/24 40.10.10.3`	Configure static route `20.20.20.0/24` by next hop interface `40.10.10.3`.
`OcNOS(config)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config)#exit`	Exit config mode.

Validation

The below shows the running configuration of the DHCPv4 Server-2 node:

```
OcNOS#show running-config dhcp
interface eth0
 ip address dhcp
!
!

ip dhcp server pool DHCP-Server-2
 network 40.10.10.0 netmask 255.255.255.0
 address range low-address 40.10.10.1 high-address 40.10.10.254
 dns-server 192.2.2.2
ip dhcp server pool DHCP-SER
 network 20.20.20.0 netmask 255.255.255.0
```

```
 address range low-address 20.20.20.1 high-address 20.20.20.30
interface xe11
 ip dhcp server
!
OcNOS#
```

DHCP Client Configuration for IPv6

`OcNOS#configure terminal`	Enter configure mode.
`OcNOS (config)#feature dhcp`	Enable the feature dhcp. This is enabled by default.
`OcNOS(config)#int xe15`	Enter interface mode `xe15`.
`OcNOS(config-if)#ipv6 address dhcp`	The client requests for the IPv6 address to the server. Once it receives the acknowledgment from the server, it assigns the IPv6 address to the interface in which this command is enabled.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.

Validation

The below shows the running configuration of the DHCPv6 client node:

```
OcNOS#show running-config dhcp
interface eth0
 ip address dhcp
!
interface xe15
 ipv6 address dhcp

OcNOS#show ipv6 int br
Interface            IPv6-Address                                    Admin-Sta
tus
cd1                  unassigned                                      [up/down]

cd3                  unassigned                                      [up/down]

ce0                  unassigned                                      [up/down]

ce2                  unassigned                                      [up/down]

eth0                 fe80::d277:ceff:fe9f:4500                       [up/up]

lo                   ::1                                             [up/up]

lo.management        ::1                                             [up/up]

xe4                  unassigned                                      [up/down]

xe5                  unassigned                                      [up/down]
```

```
xe6                 unassigned                                  [up/down]

xe7                 unassigned                                  [up/down]

xe8                 unassigned                                  [up/down]

xe9                 unassigned                                  [up/down]

xe10                unassigned                                  [up/down]

xe11                unassigned                                  [up/down]

xe12                unassigned                                  [up/down]

xe13                unassigned                                  [up/down]

xe14                unassigned                                  [up/down]

xe15                *3001::124
                    fe80::d277:ceff:feda:4511                   [up/up]

xe16                unassigned                                  [up/down]

xe17                unassigned                                  [up/down]

xe18                unassigned                                  [up/down]

xe19                unassigned                                  [up/down]

xe20                unassigned                                  [up/down]

xe21                unassigned                                  [up/down]

xe22                unassigned                                  [up/down]

xe23                unassigned                                  [up/down]

xe24                unassigned                                  [up/down]

xe25                unassigned                                  [up/down]

xe26                unassigned                                  [up/down]

xe27                unassigned                                  [up/down]

OcNOS#
OcNOS#
OcNOS#
OcNOS#
OcNOS#
```

```
OcNOS#show ipv6 int xe15 br
Interface            IPv6-Address                                    Admin-Sta
tus
xe15                 *3001::124
                     fe80::d277:ceff:feda:4511                       [up/up]
```

DHCP Relay Agent Configuration for IPv6

`OcNOS#configure terminal`	Enter configure mode.
`OcNOS(config)#ipv6 dhcp relay server-group dhcpv6-relay-gp`	Configure relay server-group group name in global mode.
`OcNOS(dhcp6-relay-group)#server 2001::2`	Configure server address `2001::2`.
`OcNOS(dhcp6-relay-group)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(dhcp6-relay-group)#exit`	Exit DHCPv6 relay group.
`OcNOS(config)#interface xe8`	Enter interface mode `xe8`.
`OcNOS(config-if)#ipv6 address 2001::3/64`	Configure IPv6 address `2001::3/64` on the interface `xe8`.
`OcNOS(config-if)#ipv6 dhcp relay uplink`	Configure relay uplink on the device connecting the server.
`OcNOS(config-if)#ipv6 dhcp relay server-select dhcpv6-relay-gp`	Configure relay server-select group name on the device connected to the server.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.
`OcNOS(config)#interface xe15`	Enter interface mode.
`OcNOS(config-if)#ipv6 address 3001::2/64`	Configure IPv6 address on the interface `xe15`.
`OcNOS(config-if)#ipv6 dhcp relay`	By default, this will be enabled. This command starts the IPv6 dhcp relay service.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.
`OcNOS(config)#ipv6 dhcp relay server-group dhcpv6-relay-gp`	Configure relay server-group group name in global mode.
`OcNOS(dhcp6-relay-group)#server 4001::2`	Configure server address `4001::2`.
`OcNOS(dhcp6-relay-group)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(dhcp6-relay-group)#exit`	Exit DHCPv6 relay group.
`OcNOS(config)#interface xe11`	Enter interface mode.
`OcNOS(config-if)#ipv6 address 4001::3/64`	Configure IPv6 `4001::3/64` address on the interface `xe11`.
`OcNOS(config-if)#ipv6 dhcp relay uplink`	Configure relay uplink on the device connecting the server.
`OcNOS(config-if)#ipv6 dhcp relay server-select dhcpv6-relay-gp`	Configure relay server-select group name on the device connected to the server.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.

Validation

The below shows the running configuration of the DHCPv6 relay agent node:

```
OcNOS#show running-config dhcp
interface eth0
 ip address dhcp
!
!

ipv6 dhcp relay server-group dhcpv6-relay-gp
 server 2001::2
 server 4001::2
interface xe8
 ipv6 dhcp relay uplink
 ipv6 dhcp relay server-select dhcpv6-relay-gp
!
interface xe11
 ipv6 dhcp relay uplink
 ipv6 dhcp relay server-select dhcpv6-relay-gp
!
interface xe15
 ipv6 dhcp relay
OcNOS#show ipv6 dhcp relay
IPv6 DHCP relay service is Enabled.
VRF Name: default
  DHCPv6 IA_PD Route injection: Disabled
  Interface               Uplink/Downlink
  ---------               -------------
  xe8                     Uplink
  xe11                    Uplink
  xe15                    Downlink
  Interface               Group-Name              Server
  ---------               -------------           ------
  xe11                    dhcpv6-relay-gp     2001::2,4001::2
OcNOS#
```

DHCP Server-1 Configuration for IPv6

`OcNOS#configure terminal`	Enter configure mode.
`OcNOS(config)#ipv6 dhcp server pool DHCPv6-Server-1`	Configure DHCP server group for server in global mode.
`OcNOS(dhcp6-config)#network 2001:: netmask 64`	Configure network `2001::` and netmask `64`.
`OcNOS(dhcp6-config)#address range low-address 2001::1 high-address 2001::124`	Configure address range from `2001::1` to `2001::124`.
`OcNOS(dhcp6-config)#ipv6 dhcp server pool DHCPv6-SER`	Configure DHCP server group for client in global mode.
`OcNOS(dhcp6-config)#network 3001:: netmask 64`	Configure network `3001::` and netmask `64`.

`OcNOS(dhcp6-config)#address range low-address 3001::1 high-address 3001::124`	Configure address range from `3001::1` to `3001::124`.
`OcNOS(dhcp6-config)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(dhcp6-config)#exit`	Exit DHCPv6 config mode.
`OcNOS(config)#interface xe8`	Enter interface mode `xe8`.
`OcNOS(config-if)#ipv6 address 2001::2/64`	Configure IPv6 address `2001::2/64` on the interface `xe8`.
`OcNOS(config-if)#ipv6 dhcp server`	Server should be configured on the interface while connected to the relay.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.
`OcNOS(config)#ipv6 route 3001::/64 2001::3`	Configure static route `3001::/64` by next hop interface `2001::3`.
`OcNOS(config)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config)#exit`	Exit config mode.

Validation

The below shows the running configuration of the DHCPv6 Server-1 node:

```
OcNOS#show running-config dhcp
interface eth0
 ip address dhcp
!
!

ipv6 dhcp server pool DHCPv6-Server-1
 network 2001:: netmask 64
 address range low-address 2001::1 high-address 2001::124
ipv6 dhcp server pool DHCPv6-SER
 network 3001:: netmask 64
 address range low-address 3001::1 high-address 3001::124
interface xe8
 ipv6 dhcp server
!
OcNOS#
```

DHCP Server-2 Configuration for IPv6

`OcNOS#configure terminal`	Enter configure mode.
`OcNOS(config)#ipv6 dhcp server pool DHCPv6-Server-2`	Configure dhcp server group for server in global mode.
`OcNOS(dhcp6-config)#network 4001:: netmask 64`	Configure network `4001::` and netmask `64`.
`OcNOS(dhcp6-config)#address range low-address 4001::1 high-address 4001::124`	Configure address range from `4001::1` to `4001::124`.

`OcNOS(dhcp6-config)#ipv6 dhcp server pool DHCPv6-SER`	Configure DHCP server group for client in global mode.
`OcNOS(dhcp6-config)#network 3001:: netmask 64`	Configure network `3001::` and netmask `64`.
`OcNOS(dhcp6-config)#address range low-address 3001::1 high-address 3001::124`	Configure address range from `3001::1` to `3001::124`.
`OcNOS(dhcp6-config)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(dhcp6-config)#exit`	Exit DHCPv6 config mode.
`OcNOS(config)#interface xe11`	Enter interface mode `xe11`.
`OcNOS(config-if)#ipv6 address 4001::2/64`	Configure IPv6 address on the interface `xe11`.
`OcNOS(config-if)#ipv6 dhcp server`	Server should be configured on the interface while connected to the relay.
`OcNOS(config-if)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config-if)#exit`	Exit interface mode.
`OcNOS(config)#ipv6 route 3001::/64 4001::3`	Configure static route `3001::/64` by next hop interface `4001::3`.
`OcNOS(config)#commit`	Commit the candidate configuration to the running configuration.
`OcNOS(config)#exit`	Exit config mode.

Validation

The below shows the running configuration of the DHCPv6 Server-2 node:

```
OcNOS#show running-config dhcp
interface eth0
 ip address dhcp
!
!

ipv6 dhcp server pool DHCPv6-Server-2
 network 4001:: netmask 64
 address range low-address 4001::1 high-address 4001::124
ipv6 dhcp server pool DHCPv6-SER
 network 3001:: netmask 64
 address range low-address 3001::1 high-address 3001::124
interface xe11
 ipv6 dhcp server
!
OcNOS#
```

New CLI Commands

ip dhcp relay server-group

Use this command to create the DHCP IPv4 server group. This group lists the servers to which DHCP Relay forwards the DHCP client requests.

Use the `no` form of this command to unconfigure the DHCP IPv4 server group.

Command Syntax

```
ip dhcp relay server-group GROUP_NAME
no ip dhcp relay server-group GROUP_NAME
```

Parameters

`GROUP_NAME`	Name of the DHCP server group (specify a maximum 63 alphanumeric characters).

Command Mode

Configure mode and VRF mode. In the configure mode, the DHCP IPv4 server group is created in the default VRF. In the configure-vrf mode, the DHCP IPv4 server group is created in the user-defined VRF.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The example below shows the creation of DHCP IPv4 server groups.

```
OcNOS(config)#ip vrf vrf1
OcNOS(config-vrf)#ip dhcp relay server-group Group1
OcNOS(dhcp-relay-group)#end
OcNOS#configure terminal
OcNOS(config)#ip dhcp relay server-group Group2
```

ip dhcp relay server-select

Use this command to attach the DHCP IPv4 server group to the DHCP relay uplink interface.

Use the `no` form of this command to remove the DHCP IPv4 server group attached to the DHCP relay interface.

Note: Attach the groups only to the DHCP relay uplink interfaces.

Command Syntax

```
ip dhcp relay server-select GROUP_NAME
no ip dhcp relay server-select
```

Parameters

`GROUP_NAME`	Name of the DHCP server group (specify a maximum 63 alphanumeric characters).

Command Mode

Interface mode.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The below example shows attaching the DHCP IPv4 server group to the DHCP relay uplink interface:

```
OcNOS#configure terminal
OcNOS(config)#interface xe1
OcNOS(config-if)#ip dhcp relay server-select group1
```

ipv6 dhcp relay server-group

Use this command to create the DHCP IPv6 server group. This group lists the servers to which DHCP relay forwards the DHCP client requests.

Use the no form of this command to unconfigure the DHCP IPv6 server group.

Command Syntax

```
ipv6 dhcp relay server-group GROUP_NAME
no ipv6 dhcp relay server-group GROUP_NAME
```

Parameters

GROUP_NAME	Name of the DHCP server group (specify a maximum of 63 alphanumeric characters).

Command Mode

Configure mode and VRF mode. In the configure mode, the DHCP IPv6 server group is created in the default VRF. In the configure-vrf mode, the DHCP IPv6 server group is created in the user-defined VRF.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The example below shows the creation of DHCP IPv6 server groups:

```
OcNOS#configure terminal
OcNOS(config)#ip vrf vrf1
OcNOS(config-vrf)#ipv6 dhcp relay server-group Group1
OcNOS(dhcp relay server-group)#end
OcNOS#configure terminal
OcNOS(config)#ipv6 dhcp relay server-group Group2
```

ipv6 dhcp relay server-select

Use this command to attach the DHCP IPv6 group to the DHCP relay uplink interface.

Use the `no` form of this command to remove the DHCP IPv6 group attached to the interface.

Note: Attach the groups only to the DHCP relay uplink interfaces.

Command Syntax

```
ipv6 dhcp relay server-select GROUP_NAME
no ipv6 dhcp relay server-select
```

Parameters

GROUP_NAME	Name of the DHCP server group (specify a maximum of 63 alphanumeric characters).

Command Mode

Interface mode.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The below example shows how to attach the DHCP IPv6 server group to the DHCP relay uplink interface:

```
#configure terminal
(config)#interface xe1
(config-if)#ipv6 dhcp relay server-select group1
```

server A.B.C.D

Use this command to add the DHCP IPv4 servers to the DHCP server group.

Use the `no` form of this command to remove the DHCP IPv4 servers from the DHCP server Group.

Note: A maximum of eight servers can be added to a DHCP group.

Command Syntax

```
server A.B.C.D
no server A.B.C.D
```

Parameters

A.B.C.D	DHCP IPv4 Relay group server address to be added in the DHCP server group.

Command Mode

DHCP Relay Group Mode.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The below example shows the addition of DHCP IPv4 servers to a DHCP server group:

```
OcNOS#configure terminal
OcNOS(config)#ip vrf vrf1
OcNOS(config-vrf)#ip dhcp relay server-group group
OcNOS(dhcp-relay-group)#server 10.12.23.205
OcNOS(dhcp-relay-group)#end
OcNOS#configure terminal
```

```
OcNOS(config)#ip dhcp relay server-group group1
OcNOS(dhcp-relay-group)#server 10.12.33.204
```

server X:X::X:X

Use this command to add the DHCP IPv6 servers to the DHCP server group.

Use the `no` form of this command to remove the DHCP IPv6 servers from the DHCP server group.

Note: A maximum of eight servers can be added to a DHCP group.

Command Syntax

```
server X:X::X:X
```

```
no server X:X::X:X
```

Parameters

`X:X::X:X`	DHCP IPv6 Relay Group server address to be added in the DHCP server group.

Command Mode

DHCPv6 Relay Group Mode.

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

The below example shows the addition of DHCP IPv6 servers to a DHCP server group:

```
OcNOS#configure terminal
OcNOS(config)#ip vrf vrf1
OcNOS(config-vrf)#ipv6 dhcp relay server-group group
OcNOS(dhcp6-relay-group)#server 2003::1
OcNOS(dhcp6-relay-group)#end

OcNOS#configure terminal
OcNOS(config)#ipv6 dhcp relay server-group group1
OcNOS(dhcp-relay-group)#server 2001::1
OcNOS(dhcp6-relay-group)end
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
DHCP	Dynamic Host Configuration Protocol
VRF	Virtual Routing and Forwarding

Glossary

The following provides definitions for key terms used throughout this document:

DHCP Client	A DHCP client is a hardware device or software that uses DHCP to get the network configuration information from a DHCP Server. VRF: VRF creates a logically isolated routing table within a single physical network device. Each VRF instance works as an independent routing instance that enables separate network traffic, maintains different routing tables, and provides network isolation.
DHCP Server	A DHCP server is a hardware device or software that leases a dynamic IP address to the DHCP client.
DHCP relay agent	A DHCP relay forwards the request from a DHCP client to the DHCP server group and takes the response from the DHCP server group to the DHCP client.
VRF	VRF creates a logically isolated routing table within a single physical network device. Each VRF instance works as an independent routing instance that enables separate network traffic, maintains different routing tables, and provides network isolation.

Improved Routing

Release 6.4.1

This section, describes the BGP peer scaling enhancements introduced in the 6.4.1 release.

- BGP Additional Path

BGP Additional Path

Overview

The Border Gateway Protocol (BGP) ADDPATH allows the advertisement of multiple paths through the same peer session for a given prefix without the new paths implicitly replacing any previous paths. This behavior promotes path diversity and reduces the severity of a network failure, thereby improving the control plane convergence in case of network failures.

Feature Characteristics

The advertisement of multiple paths in BGP is made possible by sending a BGP OPEN message to the neighbor with a BGP capability code of 69, which identifies the BGP ADD-PATH capability.

Feature	Characteristics
Address Family Identifier (AFI)	2 octets
Subsequent Address Family Identifier (SAFI)	1 octet
Send/Receive	1 octet

For a given <AFI, SAFI>, the send/receive field in the BGP TLV indicates, the sender is able to:

- Receive multiple paths from its peer (value 1).
- Send multiple paths to its peer (value 2)
- Receive and send multiple paths to its peer (value 3)
- Each alternate path is identified by a Path Identifier in addition to the address prefix

Feature	Characteristics
Path Identifier	4 octets
Length	1 octet
Prefix	variable

Benefits

This feature enables BGP add-path in the vrf address-family. In the event of a next-hop failure, BGP Add-Path improves the BGP control plane convergence time.

Prerequisites

Before configuring BGP additional paths ensure be sure of the following:

- The supported OcNOS router running a compatible release.
- Provide access to the management interface of the router.

- Understand BGP well enough to enable it BGP on an interface.

Configuration

The following sessions displays the detailed information about bgp additional paths topology, configurations, and validations.

Topology

The following topology visually represents how BGP additional paths are configured.

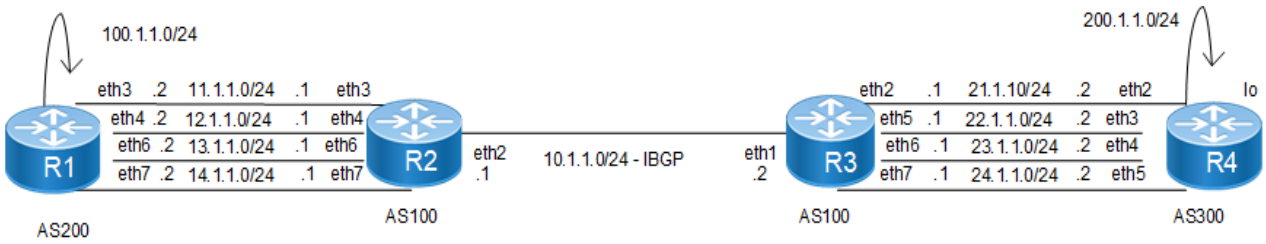

BGP Additional Path

R1

Here is the detailed configuration of router R1.

`R1#configure terminal`	Enter the Configure mode.
`R1(config)#interface eth2`	Enter the Interface mode to configure interface eth2
`R1(config-if)#ipv6 address 1001::1/64`	Configure an IPv6 address for Interface eth2
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#interface eth3`	Enter the Interface mode to configure for Interface eth3
`R1(config-if)#ipv6 address 1002::1/64`	Configure an IPv6 address for Interface eth3
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#interface eth4`	Enter the Interface mode to configure Interface eth4
`R1(config-if)#ipv6 address 1003::1/64`	Configure an IPv6 address for interface eth4
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#interface eth5`	Enter the Interface mode to configure Interface eth5
`R1(config-if)#ipv6 address 1004::1/64`	Configure an IPv6 address for Interface eth5
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#interface lo`	Enter Interface mode for loopback lo
`R1(config-if)#ipv6 address 1090::1/64`	Configure IPv6 address for Loopback interface lo
`R1(config-if)#exit`	Exit the Interface mode
`R1(config)#router bgp 200`	Enter the Router BGP mode
`R1(config-router)#neighbor 1001::2 remote-as 100`	Specify a neighbor router with a peer address and remote-as for BGP peering.
`R1(config-router)#neighbor 1002::2 remote-as 100`	Specify a neighbor router with a peer address and remote-as for BGP peering.

`R1(config-router)#neighbor 1003::2 remote-as 100`	Specify a neighbor router with a peer address and remote-as for BGP peering.
`R1(config-router)#neighbor 1004::2 remote-as 100`	Specify a neighbor router with a peer address and remote-as for BGP peering.
`R1(config-router)#address-family ipv6 unicast`	Enter address-family mode for the neighbor router session to activate.
`R1(config-router-af)#neighbor 1001::2 activate`	Activate the neighbor router with a peer address.
`R1(config-router-af)#neighbor 1002::2 activate`	Activate the neighbor router with a peer address.
`R1(config-router-af)#neighbor 1003::2 activate`	Activate the neighbor router with a peer address.
`R1(config-router-af)#neighbor 1004::2 activate`	Activate the neighbor router with a peer address.
`R1(config-router-af)#network 1090::/64`	Activate the neighbor router with a peer address.
`R1(config-router-af)#exit-address-family`	Exit the Address Family mode and return to Router mode.
`R1(config-router)#exit`	Exit the Router BGP mode and enter the Configure mode
`R1(config)#commit`	Apply commit
`R1(config)#exit`	Exit the Configure mode

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#interface eth1`	Enter Interface mode for interface eth1
`R2(config-if)#ipv6 address 3001::1/64`	Configure IPv6 address for the interface eth1
`R2(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R2(config-if)#exit`	Exit the Interface mode
`R2(config)#interface eth2`	Enter Interface mode for interface eth2
`R2(config-if)#ipv6 address 1001::2/64`	Configure IPv6 address for the Interface eth2
`R2(config-if)#exit`	Exit the interface mode
`R2(config)#interface eth3`	Enter interface mode for Interface eth3
`R2(config-if)#ipv6 address 1002::2/64`	Configure IPv6 address for the Interface eth3
`R2(config-if)#exit`	Exit the interface mode
`R2(config)#interface eth4`	Enter interface mode for Interface eth4
`R2(config-if)#ipv6 address 1003::2/64`	Configure IPv6 address for the Interface eth4
`R2(config-if)#exit`	Exit the Interface mode
`R2(config)#interface eth5`	Enter interface mode for Interface eth5
`R2(config-if)#ipv6 address 1004::2/64`	Configure IPv6 address for the interface eth5
`R2(config-if)#exit`	Exit the interface mode
`R2(config)#router bgp 100`	Enter the router bgp mode
`R2(config-router)#neighbor 3001::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.

`R2(config-router)#neighbor 1001::1 remote-as 200`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R2(config-router)#neighbor 1002::1 remote-as 200`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R2(config-router)#neighbor 1003::1 remote-as 200`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R2(config-router)#neighbor 1004::1 remote-as 200`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 1001::1 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#neighbor 1002::1 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#neighbor 1003::1 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#neighbor 1004::1 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#neighbor 3001::2 activate`	Activate the neighbor router with peer address.
`R2(config-router-af)#exit-address-family`	Exit address family mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#router ipv6 ospf`	Enter Router OSPFv3 mode.
`R2(config-router)#redistribute connected`	Configure Redistribution of Connected networks into OSPF
`R2(config-router)#exit`	Exit the router OSPF mode and enter the configure mode
`R2(config)#commit`	Apply commit
`R2(config)#exit`	Exit the configure mode

R3

Here is the detailed configuration of router R3.

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#interface eth1`	Enter Interface mode for Interface eth1
`R3(config-if)#ipv6 address 3001::2/64`	Configure IPv6 address for the Interface eth1
`R3(config-if)#ipv6 router ospf area 0`	Enable OSPFv3 routing on an interface, and assign the Area ID 0.
`R3(config-if)#exit`	Exit the Interface mode
`R3(config)#interface eth2`	Enter interface mode for Interface eth2
`R3(config-if)#ipv6 address 2001::2/64`	Configure IPv6 address for the Interface eth2
`R3(config-if)#exit`	Exit the Interface mode
`R3(config)#interface eth3`	Enter Interface mode for the Interface eth3
`R3(config-if)#ipv6 address 2002::2/64`	Configure IPv6 address for the Interface eth3
`R3(config-if)#exit`	Exit the Interface mode
`R3(config)#interface eth4`	Enter Interface mode for the Interface eth4
`R3(config-if)#ipv6 address 2003::2/64`	Configure IPv6 address for the Interface eth4
`R3(config-if)#exit`	Exit the interface mode

`R3(config)#interface eth5`	Enter Interface mode for Interface eth5
`R3(config-if)#ipv6 address 2004::2/64`	Configure IPv6 address for the Interface eth5
`R3(config-if)#exit`	Exit the Interface mode
`R3(config)#router bgp 100`	Enter the router bgp mode
`R3(config-router)#neighbor 3001::1 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R3(config-router)#neighbor 2001::1 remote-as 300`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R3(config-router)#neighbor 2002::1 remote-as 300`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R3(config-router)#neighbor 2003::1 remote-as 300`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R3(config-router)#neighbor 2004::1 remote-as 300`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R3(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R3(config-router-af)#neighbor 2001::1 activate`	Activate the neighbor router with peer address.
`R3(config-router-af)#neighbor 2002::1 activate`	Activate the neighbor router with peer address.
`R3(config-router-af)#neighbor 2003::1 activate`	Activate the neighbor router with peer address.
`R3(config-router-af)#neighbor 2004::1 activate`	Activate the neighbor router with peer address.
`R3(config-router-af)#neighbor 3001::1 activate`	Activate the neighbor router with peer address.
`R3(config-router-af)#exit-address-family`	Exit address family mode.
`R3(config-router)#exit`	Exit Router BGP mode
`R3(config)#router ipv6 ospf`	Enter Router OSPFv3 mode.
`R3(config-router)#redistribute connected`	Configure Redistribution of Connected networks into OSPF
`R3(config-router)#exit`	Exit the router OSPF mode and enter theconfigure mode
`R3(config)#commit`	Apply commit
`R3(config)#exit`	Exit the configure mode

R4

Here is the detailed configuration of router R4.

`R4#configure terminal`	Enter the Configure mode.
`R4(config)#interface eth2`	Enter interface mode for the Interface eth2
`R4(config-if)#ipv6 address 2001::1/64`	Configure IPv6 address for the Interface eth2
`R4(config-if)#exit`	Exit the Interface mode
`R4(config)#interface eth3`	Enter interface mode for Interface eth3
`R4(config-if)#ipv6 address 2002::1/64`	Configure IPv6 address for the Interface eth3
`R4(config-if)#exit`	Exit the Interface mode
`R4(config)#interface eth4`	Enter Interface mode for the Interface eth4

`R4(config-if)#ipv6 address 2003::1/64`	Configure IPv6 address for the Interface eth4
`R4(config-if)#exit`	Exit the Interface mode
`R4(config)#interface eth5`	Enter Interface mode for the Interface eth5
`R4(config-if)#ipv6 address 2004::1/64`	Configure IPv6 address for the Interface eth5
`R4(config-if)#exit`	Exit the Interface mode
`R4(config)#interface lo`	Enter interface mode for loopback lo
`R4(config-if)#ipv6 address 9999::1/64`	Configure IPv6 address for Loopback Interface lo
`R4(config-if)#exit`	Exit the interface mode
`R4(config)#router bgp 300`	Enter the router BGP mode
`R4(config-router)#neighbor 2001::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R4(config-router)#neighbor 2002::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R4(config-router)#neighbor 2003::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R4(config-router)#neighbor 2004::2 remote-as 100`	Specify a neighbor router with peer address and remote-as for BGP peering.
`R4(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R4(config-router-af)#neighbor 2001::2 activate`	Activate the neighbor router with peer address.
`R4(config-router-af)#neighbor 2002::2 activate`	Activate the neighbor router with peer address.
`R4(config-router-af)#neighbor 2003::2 activate`	Activate the neighbor router with peer address.
`R4(config-router-af)#neighbor 2004::2 activate`	Activate the neighbor router with peer address.
`R4(config-router-af)#network 9999::/64`	Activate the neighbor router with peer address.
`R4(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R4(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R4(config)#commit`	Commit the transaction
`R4(config)#exit`	Exit the configure mode

Additional Paths at the Global Level

In the following sessions additional paths at the global level is illustrated.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.

`R2(config-router-af)#bgp additional-paths send`	Configure R2 to send additional paths to all iBGP neighbors
`R2(config-router-af)#bgp additional-paths select all`	Configure R2 to select all available paths to send to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to the Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R2(config)#commit`	Commit the transaction
`R2(config)#exit`	Exit the configure mode

R3

Here is the detailed configuration of router R3.

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#router bgp 100`	Enter BGP router mode
`R3(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R3(config-router-af)#bgp additional-paths receive`	Configure R3 to receive additional paths from all iBGP neighbors
`R3(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R3(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R3(config)#commit`	Commit the transaction
`R3(config)#exit`	Exit the configure mode

Validation

The following is the validations for routers R2 and R3.

R2

The following is the validation for router.

```
#show bgp ipv6 neighbors 3001::2
BGP neighbor is 3001::2, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 10.12.5.92
  BGP state = Established, up for 00:14:55
  Last read 00:14:55, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 536 messages, 50 notifications, 0 in queue
  Sent 611 messages, 3 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 5, Offset 0, Mask 0x20
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Unicast
```

```
  BGP table version 38, neighbor version 38
  Index 5, Offset 0, Mask 0x20
  AF-dependant capabilities:
    Add-Path Send Capability : advertised
    Add-Path Receive Capability : received
  Community attribute sent to this neighbor (both)
  1 accepted prefixes
  4 announced prefixes

 Connections established 3; dropped 2
Local host: 3001::1, Local port: 38451
Foreign host: 3001::2, Foreign port: 179
Nexthop: 10.12.5.93
Nexthop global: 3001::1
Nexthop local: fe80::5054:ff:fe19:1758
BGP connection: shared network
Last Reset: 00:15:00, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)

#show bgp ipv6 summary
BGP router identifier 10.12.5.93, local AS number 100
BGP table version is 38
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
1001::1                    4   200  517        532      38       0       0
04:13:51                1
1002::1                    4   200  520        533      38       0       0
04:13:51                1
1003::1                    4   200  519        532      38       0       0
04:13:51                1
1004::1                    4   200  518        532      38       0       0
04:13:51                1
3001::2                    4   100  588        616      38       0       0
00:15:42                1

Total number of neighbors 5

Total number of Established sessions 5

#show bgp ipv6
BGP table version is 38, local router ID is 10.12.5.93
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>   1090::/64        1001::1(fe80::5054:ff:fe9c:b7e6)
                                          0        100          0       200 i
*                     1002::1(fe80::5054:ff:fe0d:f5e)
                                          0        100          0       200 i
*                     1003::1(fe80::5054:ff:fec7:1940)
                                          0        100          0       200 i
*                     1004::1(fe80::5054:ff:fe62:70d8)
                                          0        100          0       200 i
```

```
*>i  9999::/64        2001::1                0        100          0       300
i

Total number of prefixes 2

#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:54 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:54:01 2017

  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:52 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:48 2017
```

R3

The following is the validation for router R3.

```
#show bgp ipv6 neighbors 3001::1
BGP neighbor is 3001::1, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 10.12.5.93
  BGP state = Established, up for 00:29:37
  Last read 00:29:37, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
```

```
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 518 messages, 2 notifications, 0 in queue
  Sent 520 messages, 1 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 5, Offset 0, Mask 0x20
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Unicast
  BGP table version 268, neighbor version 268
  Index 1, Offset 0, Mask 0x2
  AF-dependant capabilities:
    Add-Path Send Capability : received
    Add-Path Receive Capability : advertised
  Community attribute sent to this neighbor (both)
  4 accepted prefixes
  1 announced prefixes

 Connections established 4; dropped 3
Local host: 3001::2, Local port: 179
Foreign host: 3001::1, Foreign port: 38451
Nexthop: 10.12.5.92
Nexthop global: 3001::2
Nexthop local: fe80::5054:ff:fe5d:bb79
BGP connection: shared network
Last Reset: 00:29:37, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)
#show bgp ipv6 summary
BGP router identifier 10.12.5.92, local AS number 100
BGP table version is 268
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
2001::1                    4   300  533        537     268      0      0
04:16:42               1
2002::1                    4   300  533        536     268      0      0
04:16:42               1
2003::1                    4   300  537        538     268      0      0
04:16:42               1
2004::1                    4   300  520        521     268      0      0
04:16:38               1
3001::1                    4   100  520        521     268      0      0
00:29:41               4

Total number of neighbors 5

Total number of Established sessions 5

#show bgp ipv6
BGP table version is 268, local router ID is 10.12.5.92
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>i  1090::/64        1001::1                0        100          0       200
i
* i                   1004::1                0        100          0       200
i
* i                   1003::1                0        100          0       200
i
* i                   1002::1                0        100          0       200
i
*>   9999::/64        2001::1(fe80::5054:ff:fe46:f549)
                                            0        100          0       300 i
*                     2004::1(fe80::5054:ff:feb5:9a71)
                                            0        100          0       300 i
*                     2003::1(fe80::5054:ff:fe0d:b565)
                                            0        100          0       300 i
*                     2002::1(fe80::5054:ff:fed2:4666)
                                            0        100          0       300 i

Total number of prefixes 2

R3#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 04:08:51 2017

  200
    1004::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 3      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 04:09:43 2017

  200
    1003::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 04:09:43 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 04:09:43 2017
```

Additional Paths Send and Receive at Address-family level

The following session displays the additional paths Send and Receive at Address-family level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#bgp additional-paths send-receive`	Configure R2 to send additional paths to and receive additional paths from all iBGP neighbors
`R2(config-router-af)#bgp additional-paths select all`	Configure R2 to select all available paths to send to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R2(config)#commit`	Apply commit
`R2(config)#exit`	Exit the configure mode

R3

Here is the detailed configuration of router R3.

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#router bgp 100`	Enter BGP router mode
`R3(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R3(config-router-af)#bgp additional-paths send-receive`	Configure R3 to send additional paths to and receive additional paths from all the iBGP neighbors
`R3(config-router-af)#bgp additional-paths select all`	Configure R3 to select all available paths to send to all iBGP neighbors
`R3(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R3(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R3(config)#commit`	Apply commit
`R3(config)#exit`	Exit the configure mode

Additional Paths at the Neighbor Level

The following session displays the additional paths at the neighbor level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter BGP router mode

`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 3001::2 additional-paths send-receive`	Configure R2 to send-receive additional paths to the iBGP neighbor R3
`R2(config-router-af)#neighbor 3001::2 advertise additional-paths all`	Configure R2 to advertise all available paths to the iBGP neighbor R3
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#commit`	Apply commit
`R2(config)#exit`	Exit the configure mode

R3

Here is the detailed configuration of router R3.

`R3#configure terminal`	Enter the Configure mode.
`R3(config)#router bgp 100`	Enter BGP router mode
`R3(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R3(config-router-af)#neighbor 3001::1 additional-paths send-receive`	Configure R3 to receive additional paths from the iBGP neighbor R2
`R3(config-router-af)#neighbor 3001::1 advertise additional-paths all`	Configure R2 to advertise all available paths to the iBGP neighbor R3
`R3(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R3(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R3(config)#commit`	Apply commit
`R3(config)#exit`	Exit the configure mode

Validation

The following validation for router R2 and R3 is shown below.

R2

The following validation is for router R2.

```
#show bgp ipv6 neighbors 3001::2
BGP neighbor is 3001::2, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 10.12.5.92
  BGP state = Established, up for 00:00:29
  Last read 00:00:29, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 588 messages, 51 notifications, 0 in queue
  Sent 664 messages, 4 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Minimum time between advertisement runs is 5 seconds
 For address family: IPv4 Unicast
  BGP table version 1, neighbor version 1
  Index 5, Offset 0, Mask 0x20
```

```
  Community attribute sent to this neighbor (both)
  0 accepted prefixes
  0 announced prefixes

 For address family: IPv6 Unicast
  BGP table version 64, neighbor version 64
  Index 5, Offset 0, Mask 0x20
  AF-dependant capabilities:
    Add-Path Send Capability : advertised and received
    Add-Path Receive Capability : advertised and received
  Community attribute sent to this neighbor (both)
  4 accepted prefixes
  4 announced prefixes

 Connections established 5; dropped 4
Local host: 3001::1, Local port: 179
Foreign host: 3001::2, Foreign port: 39326
Nexthop: 10.12.5.93
Nexthop global: 3001::1
Nexthop local: fe80::5054:ff:fe19:1758
BGP connection: shared network
Last Reset: 00:00:29, due to BGP Notification sent
Notification Error Message: (Cease/Other Configuration Change.)

#show bgp ipv6 summary
BGP router identifier 10.12.5.93, local AS number 100
BGP table version is 64
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                  V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
1001::1                   4   200  561        578      64       0      0
04:35:32               1
1002::1                   4   200  564        579      64       0      0
04:35:32               1
1003::1                   4   200  563        578      64       0      0
04:35:32               1
1004::1                   4   200  562        578      64       0      0
04:35:32               1
3001::2                   4   100  640        669      64       0      0
00:00:35               4

Total number of neighbors 5

Total number of Established sessions 5

#show bgp ipv6
BGP table version is 64, local router ID is 10.12.5.93
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>   1090::/64        1001::1(fe80::5054:ff:fe9c:b7e6)
                                          0        100          0       200 i
*                     1002::1(fe80::5054:ff:fe0d:f5e)
```



```
                                          0        100          0       200 i
*                     1003::1(fe80::5054:ff:fec7:1940)
                                          0        100          0       200 i
*                     1004::1(fe80::5054:ff:fe62:70d8)
                                          0        100          0       200 i
*>i  9999::/64        2001::1               0        100          0       300
i
* i                   2002::1               0        100          0       300
i
* i                   2003::1               0        100          0       300
i
* i                   2004::1               0        100          0       300
i

Total number of prefixes 2

#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:54 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:54:01 2017

  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:52 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:48 2017
```

```
#show bgp ipv6 9999::/64
BGP routing table entry for 9999::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1001::1 1002::1 1003::1 1004::1
  300
    2001::1 (metric 20) from 3001::2 (10.12.5.92)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:39 2017

  300
    2002::1 (metric 20) from 3001::2 (10.12.5.92)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: 1
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:53 2017

  300
    2003::1 (metric 20) from 3001::2 (10.12.5.92)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: 2
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:53 2017

  300
    2004::1 (metric 20) from 3001::2 (10.12.5.92)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 3      tx path_id: 3
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:53 2017
```

R3

The following validation is for router R3.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:39 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: 1
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:42 2017

  200
    1003::1 (metric 20) from 3001::1 (10.12.5.93)
```



```
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: 2
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:42 2017

  200
    1004::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 3      tx path_id: 3
      Not advertised to any peer
      Last update: Wed Jan 11 04:45:42 2017

R3#show bgp ipv6 9999::/64
BGP routing table entry for 9999::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2002::1 2003::1 2004::1
  300
    2001::1(fe80::5054:ff:fe46:f549) from 2001::1 (10.12.5.90)
    (fe80::5054:ff:fe46:f549)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::1
      Last update: Wed Jan 11 03:52:32 2017

  300
    2002::1(fe80::5054:ff:fed2:4666) from 2002::1 (10.12.5.90)
    (fe80::5054:ff:fed2:4666)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::1
      Last update: Wed Jan 11 03:52:27 2017

  300
    2003::1(fe80::5054:ff:fe0d:b565) from 2003::1 (10.12.5.90)
    (fe80::5054:ff:fe0d:b565)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::1
      Last update: Wed Jan 11 03:52:37 2017

  300
    2004::1(fe80::5054:ff:feb5:9a71) from 2004::1 (10.12.5.90)
    (fe80::5054:ff:feb5:9a71)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       3001::1
      Last update: Wed Jan 11 03:52:44 2017

#show bgp ipv6
BGP table version is 283, local router ID is 10.12.5.92
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
```

```
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>i  1090::/64        1001::1                0        100          0       200
i
* i                   1002::1                0        100          0       200
i
* i                   1003::1                0        100          0       200
i
* i                   1004::1                0        100          0       200
i
*>   9999::/64        2001::1(fe80::5054:ff:fe46:f549)
                                            0        100          0       300 i
*                     2002::1(fe80::5054:ff:fed2:4666)
                                            0        100          0       300 i
*                     2003::1(fe80::5054:ff:fe0d:b565)
                                            0        100          0       300 i
*                     2004::1(fe80::5054:ff:feb5:9a71)
                                            0        100          0       300 i

Total number of prefixes 2

#show bgp ipv6 summary
BGP router identifier 10.12.5.92, local AS number 100
BGP table version is 283
2 BGP AS-PATH entries
0 BGP community entries

Neighbor                  V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
2001::1                   4   300  556        562     282      0      0
04:28:07               1
2002::1                   4   300  556        560     283      0      0
04:28:07               1
2003::1                   4   300  560        563     282      0      0
04:28:07               1
2004::1                   4   300  543        546     283      0      0
04:28:03               1
3001::1                   4   100  551        553     283      0      0
00:04:18               4

Total number of neighbors 5

Total number of Established sessions 5

#show bgp ipv6 neighbors 3001::1
BGP neighbor is 3001::1, remote AS 100, local AS 100, internal link
  BGP version 4, remote router ID 10.12.5.93
  BGP state = Established, up for 00:05:02
  Last read 00:05:02, hold time is 90, keepalive interval is 30 seconds
  Neighbor capabilities:
    Route refresh: advertised and received (old and new)
    Address family IPv4 Unicast: advertised and received
    Address family IPv6 Unicast: advertised and received
  Received 550 messages, 3 notifications, 0 in queue
  Sent 553 messages, 2 notifications, 0 in queue
  Route refresh request: received 0, sent 0
```

```
 Minimum time between advertisement runs is 5 seconds
For address family: IPv4 Unicast
 BGP table version 1, neighbor version 1
 Index 5, Offset 0, Mask 0x20
 Community attribute sent to this neighbor (both)
 0 accepted prefixes
 0 announced prefixes

For address family: IPv6 Unicast
 BGP table version 283, neighbor version 283
 Index 1, Offset 0, Mask 0x2
 AF-dependant capabilities:
   Add-Path Send Capability : advertised and received
   Add-Path Receive Capability : advertised and received
 Community attribute sent to this neighbor (both)
 4 accepted prefixes
 4 announced prefixes

 Connections established 6; dropped 5
Local host: 3001::2, Local port: 39326
Foreign host: 3001::1, Foreign port: 179
Nexthop: 10.12.5.92
Nexthop global: 3001::2
Nexthop local: fe80::5054:ff:fe5d:bb79
BGP connection: shared network
Last Reset: 00:05:07, due to BGP Notification received
Notification Error Message: (Cease/Other Configuration Change.)
```

Selection of all Additional Paths at the Address-family Level

The following are the configurations and validations for additional paths at the address-family level.

R2

Here is the detailed configuration of router R2.

`#configure terminal`	Enter the Configure mode.
`(config)#router bgp 100`	Enter BGP router mode
`(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`(config-router-af)#bgp additional-paths send-receive`	Configure R2 to send additional paths to and receive additional paths from all iBGP neighbors
`(config-router-af)#bgp additional-paths select all`	Configure R2 to select all available paths to send to all iBGP neighbors
`(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`(config)#commit`	Apply the commit
`(config)#exit`	Exit the configure mode

Selection of all Additional Paths at the Neighbor Level

The following are the configurations and validations for additional paths at the neighbor level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to and receive additional paths from all iBGP neighbors
`R2(config-router-af)#neighbor 3001::2 advertise additional-paths all`	Configure R2 to select all available paths to send to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Validation

The following is the validations for routers R2 and R3.

R2

The following is the validation for router R2.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:54 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:54:01 2017
```

```
  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:52 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 3
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 03:53:48 2017
```

R3

The following is the validation for router R3.

```
#show bgp ipv6 1090::
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 05:52:01 2017

  200
    1004::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 3      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 05:52:43 2017

  200
    1003::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 05:52:43 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 05:52:43 2017
```

Selection of Best 2 Additional Paths at AF Level

The following are the configurations and validations for best 2 additional paths at AF level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#bgp additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`R2(config-router-af)#bgp additional-paths select best 2`	Configure R2 to select best 2 out of all available paths to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Selection of Best 2 Additional Paths at the Neighbor Level

The following are the configurations and validations for best 2 additional paths at neighbor level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`R2(config-router-af)#neighbor 3001::2 advertise additional-paths best 2`	Configure R2 to advertise best 2 out of all available paths to R3
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the config mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Validation

The following is the validations for routers R2 and R3.

R2

The following is the validation for router R2.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017

  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:34:49 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:34:49 2017
```

R3

The following is the validation for router R3.

```
#show bgp ipv6 1090::
BGP routing table entry for 1090::/64
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
```

```
      Last update: Wed Jan 11 06:34:49 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:34:49 2017
#show bgp ipv6
BGP table version is 407, local router ID is 10.12.5.92
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>i  1090::/64        1001::1               0        100          0       200
i
* i                   1002::1               0        100          0       200
i
*>   9999::/64        2001::1(fe80::5054:ff:fe46:f549)
                                           0        100          0       300 i
*                     2002::1(fe80::5054:ff:fed2:4666)
                                           0        100          0       300 i
*                     2003::1(fe80::5054:ff:fe0d:b565)
                                           0        100          0       300 i
*                     2004::1(fe80::5054:ff:feb5:9a71)
                                           0        100          0       300 i

Total number of prefixes 2
```

Selection of Best 3 Additional Paths at the AF Level

The following are the configurations and validations for best 3 additional paths at the AF level.

R2

THere is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#bgp additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`R2(config-router-af)#bgp additional-paths select best 3`	Configure R2 to select best 3 out of all available paths to all iBGP neighbors
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Selection of Best 3 Additional Paths at the Neighbor Level

The following are the configurations and validations for best 3 additional paths at neighbor level.

R2

Here is the detailed configuration of router R2.

`R2#configure terminal`	Enter the Configure mode.
`R2(config)#router bgp 100`	Enter the BGP router mode
`R2(config-router)#address-family ipv6 unicast`	Enter address-family mode for neighbor router session to activate.
`R2(config-router-af)#neighbor 3001::2 additional-paths send`	Configure R2 to send additional paths to the iBGP neighbor R3
`R2(config-router-af)#neighbor 3001::2 advertise additional-paths best 3`	Configure R2 to advertise best 3 out of all available paths to R3
`R2(config-router-af)#exit-address-family`	Exit Address Family mode and return to Router mode.
`R2(config-router)#exit`	Exit the router BGP mode and enter the configure mode
`R2(config)#commit`	Apply the commit
`R2(config)#exit`	Exit the configure mode

Validation

The following is the validations for routers R2 and R3.

R2

The following is the validation for router R2.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (4 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  1002::1 1003::1 1004::1
  200
    1001::1(fe80::5054:ff:fe9c:b7e6) from 1001::1 (10.12.5.144)
    (fe80::5054:ff:fe9c:b7e6)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx path_id: -1      tx path_id: 0
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017

  200
    1002::1(fe80::5054:ff:fe0d:f5e) from 1002::1 (10.12.5.144)
    (fe80::5054:ff:fe0d:f5e)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 1
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017
```

```
  200
    1003::1(fe80::5054:ff:fec7:1940) from 1003::1 (10.12.5.144)
    (fe80::5054:ff:fec7:1940)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: 2
      Advertised to non peer-group peers:
       3001::2
      Last update: Wed Jan 11 06:34:49 2017

  200
    1004::1(fe80::5054:ff:fe62:70d8) from 1004::1 (10.12.5.144)
    (fe80::5054:ff:fe62:70d8)
      Origin IGP, metric 0, localpref 100, valid, external
      rx path_id: -1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:34:49 2017
```

R3

The following is the validation for router R3.

```
#show bgp ipv6 1090::/64
BGP routing table entry for 1090::/64
Paths: (3 available, best #1, table Default-IP-Routing-Table)
  Advertised to non peer-group peers:
  2001::1 2002::1 2003::1 2004::1
  200
    1001::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      rx path_id: 0      tx path_id: 0
      Not advertised to any peer
      Last update: Wed Jan 11 06:36:11 2017

  200
    1003::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 2      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:36:53 2017

  200
    1002::1 (metric 20) from 3001::1 (10.12.5.93)
      Origin IGP, metric 0, localpref 100, valid, internal
      rx path_id: 1      tx path_id: -1
      Not advertised to any peer
      Last update: Wed Jan 11 06:36:53 2017

#show bgp ipv6
BGP table version is 410, local router ID is 10.12.5.92
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

    Network          Next Hop            Metric    LocPrf Weight Path
*>i  1090::/64        1001::1               0        100          0       200
i
```



```
* i                    1003::1                0        100          0       200
i
* i                    1002::1                0        100          0       200
i
*>   9999::/64         2001::1(fe80::5054:ff:fe46:f549)
                                              0        100          0       300 i
*                      2002::1(fe80::5054:ff:fed2:4666)
                                              0        100          0       300 i
*                      2003::1(fe80::5054:ff:fe0d:b565)
                                              0        100          0       300 i
*                      2004::1(fe80::5054:ff:feb5:9a71)
                                              0        100          0       300 i

Total number of prefixes 2
```

Implementation Examples

The following examples shows the bgp additional path entries.

```
OcNOS#show ip bg summary
BGP router identifier 33.33.33.1, local AS number 400
BGP table version is 2
1 BGP AS-PATH entries
0 BGP community entries
Neighbor                 V   AS   MsgRcv    MsgSen TblVer   InQ   OutQ    Up/
Down   State/PfxRcd
1.1.1.2                  4   100   19         17       2      0      0
00:06:35               0
Total number of neighbors 1
Total number of Established sessions 1
```

CLI Commands

The BGP additional path introduces the following configuration commands.

bgp additional-paths send-receive

Use this command is to enable BGP additional paths send and receive global mode commands in ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path send and receive global mode commands in ipv4 vrf address-family.

Command Syntax

```
bgp additional-paths (send-receive)
no bgp additional-paths (send-receive)
```

Parameters

`send`	Send additional paths to neighbors.
`receive`	Receive additional paths from neighbors.
`send-receive`	Send and receive additional paths from neighbors.

`select`	Selection criteria to pick the paths.

Default

By default, BGP additional paths is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#bgp additional-paths send
(config-router-af)#no bgp additional-paths send
```

bgp additional-paths select best 3

Use this command is to enable BGP additional best 3 paths in global mode ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path select best 3 paths in global mode ipv4 vrf address-family.

Command Syntax

```
bgp additional-paths(select best 3)
no bgp additional-paths(select best 3)
```

Parameters

`best`	Select best N paths.
`3`	Number of best paths in additional paths to be selected.
`select`	Selection criteria to pick the paths.

Default

By default, BGP additional paths is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
```

```
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#bgp additional-paths send
(config-router-af)#no bgp additional-paths send
```

bgp additional-paths select all

Use this command to enable BGP additional paths select all in global mode commands in ipv4 vrf address-family.

Use the `no` parameter with this command to disable all selected BGP add-paths in global mode commands in ipv4 vrf address-family..

Command Syntax

```
bgp additional-paths(select all)
no bgp additional-paths(select all)
```

Parameters

`all`	Select all available paths.
`select`	Selection criteria to pick the paths.

Default

By default, BGP additional paths is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#bgp additional-paths send
(config-router-af)#no bgp additional-paths send
```

neighbor A.B.C.D additional-paths send | receive | send-receive

Use this command to enable BGP add-path at neighbor level to send and receive neighbor level commands added in ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path at neighbor level to send and receive neighbor level commands added in ipv4 vrf address-family.

Command Syntax

```
neighbor A.B.C.D additional-paths (send|receive|send-receive|)
no neighbor A.B.C.D additional-paths (send|receive|send-receive|)
```

Parameters

`send`	Send additional paths to neighbors.
`receive`	Receive additional paths from neighbors.
`send-receive`	Send and receive additional paths from neighbors.

Default

By default, neighbor advertise additional path is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#neighbor  1.1.1.2 advertise additional-paths all
(config-router-af)#no neighbor  1.1.1.2 advertise additional-paths all
```

neighbor A.B.C.D additional-paths all

Use this command to enable BGP add-path at all neighbor level commands added in ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path at neighbor level commands added in ipv4 vrf address-family.

Command Syntax

```
neighbor A.B.C.D additional-paths all
no neighbor A.B.C.D additional-paths all
```

Parameters

`all`	Select all available paths

Default

By default, neighbor advertise additional path is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
```

```
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#neighbor  1.1.1.2 advertise additional-paths all
(config-router-af)#no neighbor  1.1.1.2 advertise additional-paths all
```

neighbor A.B.C.D additional-paths best <2-3>

Use this command to enable BGP add-path at all neighbor level commands added in ipv4 vrf address-family.

Use the `no` parameter with this command to disable BGP add-path at neighbor level commands added in ipv4 vrf address-family.

Command Syntax

```
neighbor A.B.C.D additional-paths all
no neighbor A.B.C.D additional-paths all
```

Parameters

`best`	Select best N paths.
`<2-3>`	Number of best paths in additional paths to be selected.

Default

By default, neighbor advertise additional path is disabled.

Command Mode

Address Family IPV4 VRF Mode.

Applicability

This command was introduced before OcNOS version 6.4.1.

Examples

```
OcNOS#configure terminal
(config)#router bgp 2
(config-router)#address-family ipv4 vrf ip1
(config-router-af)#neighbor  1.1.1.2 advertise additional-paths all
(config-router-af)#no neighbor  1.1.1.2 advertise additional-paths all
```

Troubleshooting

BGP additional paths for DC is a new feature when it is configured neighbour level resets only the particular peer and global level resets all the peers.

Abbreviations

List key terms used in this document and add the term and explanation to our existing Glossary.

Acronym	Description
BGP	Border Gateway Protocol
CLI	Command Line Interface
TLV	Type Length Values

Improved Network Resilience

Release 6.4.1

This section, describes the network fail over and error handling enhancements introduced in the 6.4.1 release.

- RSVP Detour Over Ring Topology
- Commit Rollback

RSVP Detour Over Ring Topology

Overview

In OcNOS, this feature allows the detour formation in the ring topology to enhance the routing experience. The detour formation is a local protection mechanism to reroute the data traffic when a failure or congestion occurs in the primary Label Switched Path (LSP). In Multiprotocol Label Switching (MPLS), the primary LSP is the default path through which the data travels from the source to the destination node.

Feature Characteristics

This feature allows detour to take the upstream path of protected LSP, allowing a detour based protection in a ring topology. The upstream path of the protected LSP is the section of the network that precedes the PLR node in the network. This feature works for both path and sender-template method of detour formation. For the inter-op solutions that do not support the sender-template method, use the path method of detour formation.

In the below diagram, the data traffic path highlighted in green dots is the primary LSP. The link shown with the red cross is locally protected at the Point of Local Repair (PLR) node. A PLR node is a network device that reacts and takes action when a link fails. For continued data traffic flow, detour occurs through the red dotted line. Detour in MPLS is an alternate path used when the primary LSP encounters disruption or congestion.

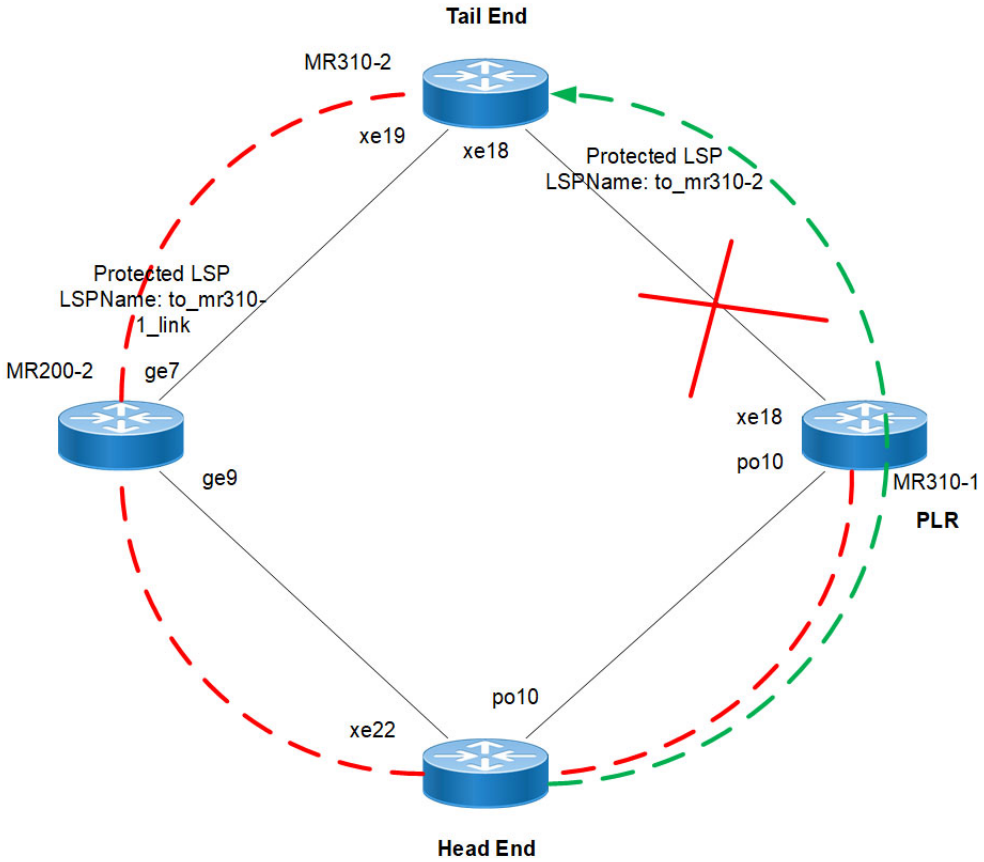

RSVP-TE FRR failover ring topology Feature Characteristics

Benefits

This feature helps detour the data traffic when there is a link or node failure, keeping the data traffic loss to a minimum (less than 50ms when BFD negotiated for fastest detection).

Prerequisite

Before the detour configuration in a ring topology, configure the RSVP tunnel with fast reroute protection of the one-to-one method.

For more information, refer to the Fast Reroute Configuration (one-to-one method) section of the RSVP Detour Over Ring Topology chapter in the *OcNOS Multi-Protocol Label Switching Guide,* Release 6.4.1.

Configuration

This section shows the configuration procedure to create a detour in the ring topology.

Topology

Configure the primary LSP in the below ring topology from the head end to the tail end.

For example, consider PE1 as the head end and PE2 as the tail end, and the primary LSP is via R1, R2, and R3. In this case, first configure the Fast Reroute Configuration (one-to-one method) on the PE1 and PE2 and then configure the detour-allow-primary-upstream-path command in all the nodes. For example, if the link between R3 and PE2 is down, the detour follows via primary LSP to reach PE2.

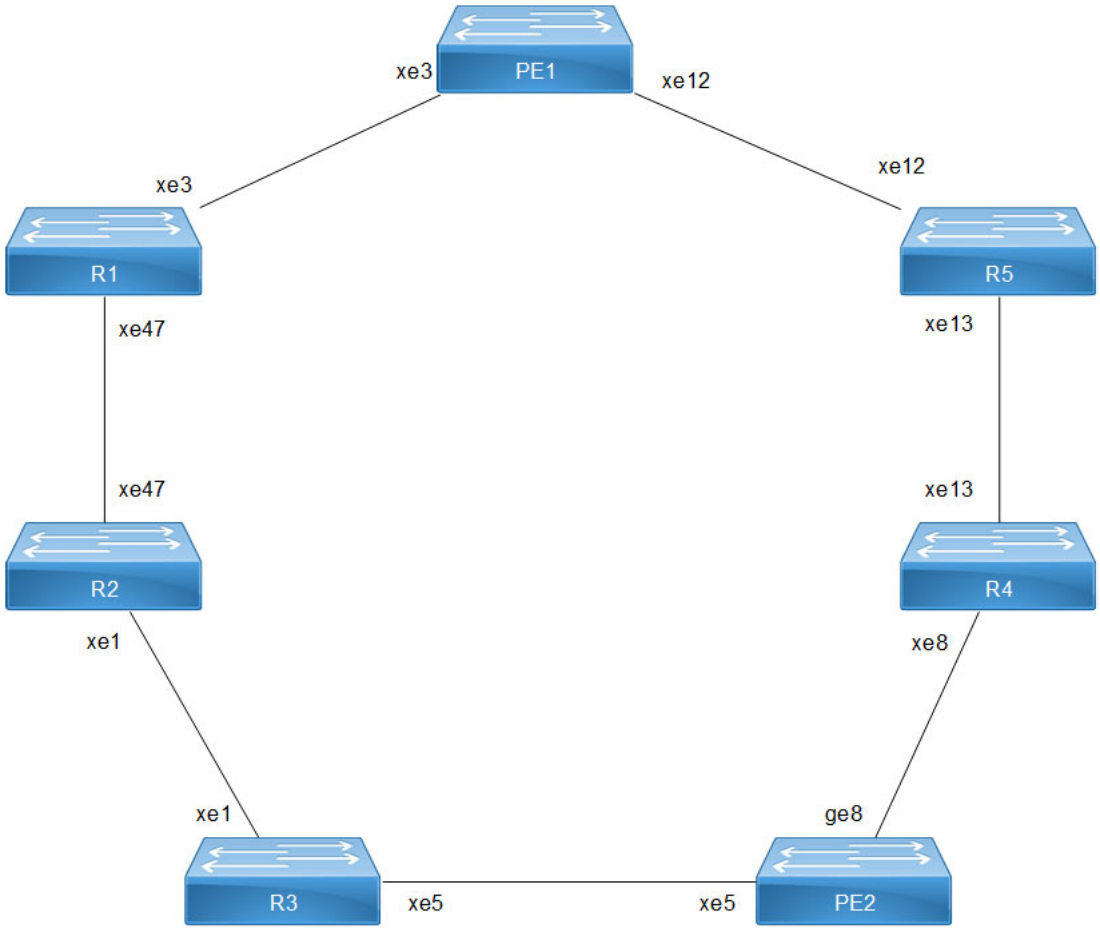

RSVP-TE FRR failover ring topology - 1:1 Detour

PE1 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

`PE1#configure terminal`	Enter configure mode.
`PE1(config)#interface xe3`	Enter interface mode `xe3`.
`PE1(config-if)#ip address 61.61.61.3/24`	Configure IPv4 address `61.61.61.3.24`.
`PE1(config-if)#label-switching`	Configure label switching on `xe3`.
`PE1(config-if)#enable-rsvp`	Enable RSVP on `xe3`.
`PE1(config-if)#exit`	Exit interface mode.
`PE1(config)#interface xe12`	Enter interface mode `xe12`.
`PE1(config-if)#ip address 58.58.58.2/24`	Configure IPv4 address `58.58.58.2/24`.
`PE1(config-if)#label-switching`	Configure label switching on `xe12`.
`PE1(config-if)#enable-rsvp`	Enable RSVP on `xe12`.
`PE1(config-if)#exit`	Exit interface mode.
`PE1(config)#interface lo`	Enter loopback interface mode.

`PE1(config-if)#ip address 26.26.26.26/32 secondary`	Configure IPv4 address `26.26.26.26/32`.
`PE1(config-if)#exit`	Exit interface mode.
`PE1(config)#router ospf 100`	Enter OSPF router mode.
`PE1(config-router)#ospf router-id 26.26.26.26`	Assign router ID `26.26.26.26` for OSPF.
`PE1(config-router)#network 26.26.26.26/32 area 0.0.0.0`	Define network `26.26.26.26/32` under router OSPF.
`PE1(config-router)#network 58.58.58.0/24 area 0.0.0.0`	Define network `58.58.58.0/24` under router OSPF.
`PE1(config-router)#network 61.61.61.0/24 area 0.0.0.0`	Define network `61.61.61.0/24` under router OSPF.
`PE1(config-router)#exit`	Exit router OSPF mode.
`PE1(config)#commit`	Commit the transaction.
`PE1(config)#exit`	Exit the configure mode.

PE1 - RSVP Configurations

This section shows:

1. The configuration of detour to take the upstream path of protected LSP.
2. The configuration of the primary LSP and attaching it to the RSVP trunk.
3. The configuration of the FRR.

`PE1#configure terminal`	Enter configure mode.
`PE1(config)#router rsvp`	Enable RSVP globally.
`PE1(config-router)#detour-allow-primary-upstream-path`	Configure this CLI to allow detour to take primary upstream path.
`PE1(config-router)#exit`	Exit router RSVP mode.
`PE1(config)#rsvp-path PE1-PE2-01 mpls`	Configure RSVP path PE1-PE2-01 and enter path mode.
`PE1(config-path)#61.61.61.2 strict`	Configure this explicit route path as a strict hop.
`PE1(config-path)#23.23.23.3 strict`	Configure this explicit route path as a strict hop.
`PE1(config-path)#41.41.41.3 strict`	Configure this explicit route path as a strict hop.
`PE1(config-path)#56.56.56.3 strict`	Configure this explicit route path as a strict hop.
`PE1(config-path)#rsvp-trunk TR-PE1-PE2-MP-01 ipv4`	Create an RSVP trunk TR-PE1-PE2-MP-01 and enter the trunk mode.
`PE1(config-trunk)#primary fast-reroute protection one-to-one`	Configure primary fast reroute protection.
`PE1(config-trunk)#primary fast-reroute node-protection`	Configure node protection.
`PE1(config-trunk)#primary path PE1-PE2-01`	Configure trunk PE1-PE2-01 to use as the primary LSP.
`PE1(config-trunk)#from 26.26.26.26`	Assign the source loopback address `26.26.26.26` to the RSVP trunk.
`PE1(config-trunk)#to 22.22.22.22`	Assign the destination loopback address `22.22.22.22` to the RSVP trunk.
`PE1(config-trunk)#exit`	Exit router RSVP trunk mode.

`PE1(config)#commit`	Commit the transaction.
`PE1(config)#exit`	Exit the configure mode.

R1 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

`R1#configure terminal`	Enter configure mode.
`R1(config)#interface xe3`	Enter interface mode `xe3`.
`R1(config-if)#ip address 61.61.61.2/24`	Configure IPv4 address `61.61.61.2/24`.
`R1(config-if)#label-switching`	Configure label switching on `xe3`.
`R1(config-if)#enable-rsvp`	Enable RSVP on interface `xe3`.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface xe47`	Enter interface mode `xe47`.
`R1(config-if)#ip address 23.23.23.2/24`	Configure IPv4 address `23.23.23.2/24`.
`R1(config-if)#label-switching`	Configure label switching on `xe47`.
`R1(config-if)#enable-rsvp`	Enable RSVP on interface `xe47`.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#interface lo`	Enter loopback interface mode.
`R1(config-if)#ip address 24.24.24.24/32 secondary`	Configure IPv4 address `24.24.24.24/32`.
`R1(config-if)#exit`	Exit interface mode.
`R1(config)#router ospf 100`	Enter OSPF router mode.
`R1(config-router)#ospf router-id 24.24.24.24`	Assign router-id for OSPF.
`R1(config-router)#network 23.23.23.0/24 area 0.0.0.0`	Define network `23.23.23.0/24` under router OSPF.
`R1(config-router)#network 24.24.24.24/32 area 0.0.0.0`	Define network `24.24.24.24/32` under router OSPF.
`R1(config-router)#network 61.61.61.0/24 area 0.0.0.0`	Define network `61.61.61.0/24` under router OSPF.
`R1(config-router)#exit`	Exit router OSPF mode.
`R1(config)#commit`	Commit the transaction.
`R1(config)#exit`	Exit the configure mode.

R1 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

`R1#configure terminal`	Enter configure mode.
`R1(config)#router rsvp`	Enable RSVP globally.
`R1(config-router)#detour-allow-primary-upstream-path`	Configure this CLI to allow detour to take primary upstream path.
`R1(config-router)#exit`	Exit router RSVP mode.
`R1(config)#commit`	Commit the transaction.
`R1(config)#exit`	Exit the configure mode.

R2 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

`R2#configure terminal`	Enter configure mode.
`R2(config)#interface xe1`	Enter interface mode `xe1`.
`R2(config-if)#ip address 41.41.41.2/24`	Configure IPv4 address `41.41.41.2/24`.
`R2(config-if)#label-switching`	Configure label switching on `xe1`.
`R2(config-if)#enable-rsvp`	Enable RSVP on `xe1`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface xe47`	Enter interface mode `xe47`.
`R2(config-if)#ip address 23.23.23.3/24`	Configure IPv4 address `23.23.23.3/24`.
`R2(config-if)#label-switching`	Configure label switching on `xe47`.
`R2(config-if)#enable-rsvp`	Enable RSVP on `xe47`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#interface lo`	Enter loopback interface mode.
`R2(config-if)#ip address 88.88.88.88/32 secondary`	Configure IPv4 address `88.88.88.88/32`.
`R2(config-if)#exit`	Exit interface mode.
`R2(config)#router ospf 100`	Enter OSPF router mode.
`R2(config-router)#ospf router-id 88.88.88.88`	Assign router-id `88.88.88.88` for OSPF.
`R2(config-router)#network 23.23.23.0/24 area 0.0.0.0`	Define network `23.23.23.0/24` under router OSPF.
`R2(config-router)#network 41.41.41.0/24 area 0.0.0.0`	Define network `41.41.41.0/24` under router OSPF.
`R2(config-router)#network 88.88.88.88/32 area 0.0.0.0`	Define network `88.88.88.88/32` under router OSPF.
`R2(config-router)#exit`	Exit router OSPF mode.
`R2(config)#commit`	Commit the transaction.
`R2(config)#exit`	Exit the configure mode.

R2 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

`R2#configure terminal`	Enter configure mode.
`R2(config)#router rsvp`	Enable RSVP globally.
`R2(config-router)#detour-allow-primary-upstream-path`	Configure this CLI to allow detour to take primary upstream path.
`R2(config-router)#exit`	Exit router RSVP mode.
`R2(config)#commit`	Commit the transaction.
`R2(config)#exit`	Exit the configure mode.

R3 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

`R3#configure terminal`	Enter configure mode.
`R3(config)#interface xe1`	Enter interface mode `xe1`.
`R3(config-if)#ip address 41.41.41.3/24`	Configure IPv4 address `41.41.41.3/24`.
`R3(config-if)#label-switching`	Configure label switching on `xe1`.
`R3(config-if)#enable-rsvp`	Enable RSVP on `xe1`.
`R3(config-if)#exit`	Exit interface mode.
`R3(config)#interface xe5`	Enter interface mode `xe5`.
`R3(config-if)#ip address 56.56.56.2/24`	Configure IPv4 address `56.56.56.2/24`.
`R3(config-if)#label-switching`	Configure label switching on `xe5`.
`R3(config-if)#enable-rsvp`	Enable RSVP on `xe5`.
`R3(config-if)#exit`	Exit interface mode.
`R3(config)#interface lo`	Enter loopback interface mode.
`R3(config-if)#ip address 99.99.99.99/32 secondary`	Configure IPv4 address `99.99.99.99/32`.
`R3(config-if)#exit`	Exit interface mode.
`R3(config)#router ospf 100`	Enter OSPF router mode.
`R3(config-router)#ospf router-id 99.99.99.99`	Assign router-id for OSPF.
`R3(config-router)#network 41.41.41.0/24 area 0.0.0.0`	Define network `41.41.41.0/24` under router OSPF.
`R3(config-router)#network 56.56.56.0/24 area 0.0.0.0`	Define network `56.56.56.0/24` under router OSPF.
`R3(config-router)#network 99.99.99.99/32 area 0.0.0.0`	Define network `99.99.99.99/32` under router OSPF.
`R3(config-router)#exit`	Exit router OSPF mode.
`R3(config)#commit`	Commit the transaction.
`R3(config)#exit`	Exit the configure mode.

R3 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

`R3#configure terminal`	Enter configure mode.
`R3(config)#router rsvp`	Enable RSVP globally.
`R3(config-router)#detour-allow-primary-upstream-path`	Configure this CLI to allow detour to take primary upstream path.
`R3(config-router)#exit`	Exit router RSVP mode.
`R3(config)#commit`	Commit the transaction.
`R3(config)#exit`	Exit the configure mode.

R5 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

`R5#configure terminal`	Enter configure mode.
`R5(config)#interface xe1`	Enter interface mode `58.58.58.3/24`.
`R5(config-if)#ip address 58.58.58.3/24`	Configure IPv4 address.
`R5(config-if)#label-switching`	Configure label switching on `xe1`.
`R5(config-if)#enable-rsvp`	Enable RSVP on `xe1`.
`R5(config-if)#exit`	Exit interface mode.
`R5(config)#interface xe13`	Enter interface mode `xe13`.
`R5(config-if)#ip address 54.54.54.4/24`	Configure IPv4 address `54.54.54.4/24`.
`R5(config-if)#label-switching`	Configure label switching on `xe13`.
`R5(config-if)#enable-rsvp`	Enable RSVP on `xe13`.
`R5(config-if)#exit`	Exit interface mode.
`R5(config)#interface lo`	Enter loopback interface mode.
`R5(config-if)#ip address 17.17.17.17/32 secondary`	Configure IPv4 address `17.17.17.17/32`.
`R5(config-if)#exit`	Exit interface mode.
`R5(config)#router ospf 100`	Enter OSPF router mode.
`R5(config-router)#ospf router-id 17.17.17.17`	Assign router-id for OSPF.
`R5(config-router)#network 17.17.17.17/32 area 0.0.0.0`	Define network `17.17.17.17/32` under router OSPF.
`R5(config-router)#network 54.54.54.0/24 area 0.0.0.0`	Define network `54.54.54.0/24` under router OSPF.
`R5(config-router)#network 58.58.58.0/24 area 0.0.0.0`	Define network `58.58.58.0/24` under router OSPF.
`R5(config-router)#exit`	Exit router OSPF mode.
`R5(config)#commit`	Commit the transaction.
`R5(config)#exit`	Exit the configure mode.

R5 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

`R5#configure terminal`	Enter configure mode.
`R5(config)#router rsvp`	Enable RSVP globally.
`R5(config-router)#detour-allow-primary-upstream-path`	Configure this CLI to allow detour to take primary upstream path
`R5(config-router)#exit`	Exit router RSVP mode
`R5(config)#commit`	Commit the transaction.
`R5(config)#exit`	Exit the configure mode.

R4 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

`R4#configure terminal`	Enter configure mode.
`R4(config)#interface xe13`	Enter interface mode `xe13`.
`R4(config-if)#ip address 54.54.54.3/24`	Configure IPv4 address `54.54.54.3/24`.
`R4(config-if)#label-switching`	Configure label switching on `xe13`.
`R4(config-if)#enable-rsvp`	Enable RSVP on interface `xe13`.
`R4(config-if)#exit`	Exit interface mode.
`R4(config)#interface xe8`	Enter interface mode `xe8`.
`R4(config-if)#ip address 62.62.62.3/24`	Configure IPv4 address `62.62.62.3/24`.
`R4(config-if)#label-switching`	Configure label switching on `xe8`.
`R4(config-if)#enable-rsvp`	Enable RSVP on `xe8`.
`R4(config-if)#exit`	Exit interface mode.
`R4(config)#interface lo`	Enter loopback interface mode.
`R4(config-if)#ip address 48.48.48.48/32 secondary`	Configure IPv4 address `48.48.48.48/32`.
`R4(config-if)#exit`	Exit interface mode.
`R4(config)#router ospf 100`	Enter OSPF router mode.
`R4(config-router)#ospf router-id 48.48.48.48`	Assign router-id for OSPF.
`R4(config-router)#network 48.48.48.48/32 area 0.0.0.0`	Define network `48.48.48.48/32` under router OSPF.
`R4(config-router)#network 54.54.54.0/24 area 0.0.0.0`	Define network `54.54.54.0/24` under router OSPF.
`R4(config-router)#network 62.62.62.0/24 area 0.0.0.0`	Define network 62.62.62.0/24 under router OSPF.
`R4(config-router)#exit`	Exit router OSPF mode.
`R4(config)#commit`	Commit the transaction.
`R4(config)#exit`	Exit the configure mode.

R4 - RSVP Configurations

This section shows how to configure the detour to take the upstream path of protected LSP.

`R4#configure terminal`	Enter configure mode.
`R4(config)#router rsvp`	Enable RSVP globally.
`R4(config-router)#detour-allow-primary-upstream-path`	Configure this CLI to allow detour to take primary upstream path.
`R4(config-router)#exit`	Exit router RSVP mode.
`R4(config)#commit`	Commit the transaction.
`R4(config)#exit`	Exit the configure mode.

PE2 - OSPF Configurations

This section shows how to configure the Open Shortest Path First (OSPF) as Interior Gateway Protocol (IGP).

`PE2#configure terminal`	Enter configure mode.
`PE2(config)#interface xe5`	Enter interface mode `xe5`.
`PE2(config-if)#ip address 56.56.56.3/24`	Configure IPv4 address `56.56.56.3/24`.
`PE2(config-if)#label-switching`	Configure label switching on `xe5`.
`PE2(config-if)#enable-rsvp`	Enable RSVP on `xe5`.
`PE2(config-if)#exit`	Exit interface mode.
`PE2(config)#interface ge8`	Enter interface mode `ge8`.
`PE2(config-if)#ip address 62.62.62.2/24`	Configure IPv4 address `62.62.62.2/24`.
`PE2(config-if)#label-switching`	Configure label switching on `ge8`.
`PE2(config-if)#enable-rsvp`	Enable RSVP on `ge8`.
`PE2(config-if)#exit`	Exit interface mode.
`PE2(config)#interface lo`	Enter loopback interface mode.
`PE2(config-if)#ip address 22.22.22.22/32 secondary`	Configure IPv4 address `22.22.22.22/32`.
`PE2(config-if)#exit`	Exit interface mode.
`PE2(config)#router ospf 100`	Enter OSPF router mode.
`PE2(config-router)#ospf router-id 22.22.22.22`	Assign router-id for OSPF.
`PE2(config-router)#network 22.22.22.22/32 area 0.0.0.0`	Define network `22.22.22.22/32` under router OSPF.
`PE2(config-router)#network 56.56.56.0/24 area 0.0.0.0`	Define network `56.56.56.0/24` under router OSPF.
`PE2(config-router)#network 62.62.62.0/24 area 0.0.0.0`	Define network `62.62.62.0/24` under router OSPF.
`PE2(config-router)#exit`	Exit router OSPF mode.
`PE2(config)#commit`	Commit the transaction.
`PE2(config)#exit`	Exit the configure mode.

PE2 - RSVP Configurations

This section shows:

1. The configuration of detour to take the upstream path of protected LSP.
2. The configuration of the primary LSP and attaching it to the RSVP trunk.
3. The configuration of the FRR.

`PE2#configure terminal`	Enter configure mode.
`PE2(config)#router rsvp`	Enable RSVP globally.
`PE2(config-router)#detour-allow-primary-upstream-path`	Configure this CLI to allow detour to take primary upstream path.
`PE2(config-router)#exit`	Exit router RSVP mode.
`PE2(config)#rsvp-path PE2-PE1-01 mpls`	Configure RSVP path PE2-PE1-01 and enter path mode.
`PE2(config-path)#56.56.56.2 strict`	Configure this explicit route path as a strict hop.
`PE2(config-path)#41.41.41.2 strict`	Configure this explicit route path as a strict hop.

`PE2(config-path)#23.23.23.2 strict`	Configure this explicit route path as a strict hop.
`PE2(config-path)#61.61.61.3 strict`	Configure this explicit route path as a strict hop.
`PE2(config-router)#exit`	Exit path mode.
`PE2(config-path)#rsvp-trunk TR-PE2-PE1-MP-01 ipv4`	Create an RSVP trunk TR-PE2-PE1-MP-01 and enter the Trunk mode.
`PE2(config-trunk)#primary fast-reroute protection one-to-one`	Configure primary fast-reroute protection.
`PE2(config-trunk)#primary fast-reroute node-protection`	Configure node protection.
`PE2(config-trunk)#primary path PE2-PE1-01`	Configure trunk PE2-PE1-01 to use as the primary LSP.
`PE2(config-trunk)#from 22.22.22.22`	Assign the source loopback address `22.22.22.22` to the RSVP trunk.
`PE2(config-trunk)#to 26.26.26.26`	Assign the destination loopback address `26.26.26.26` to the RSVP trunk.
`PE2(config-trunk)#exit`	Exit router RSVP trunk mode.
`PE2(config)#commit`	Commit the transaction.
`PE2(config)#exit`	Exit the configure mode.

Validation

PE1

Below is the validation output of RSVP LSPs from PE1 to PE2 via R1>R2>R3:

```
#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     26.26.26.26     5001     2205     PRI    TR-PE1-PE2-MP-01-Primary        UP    02:12:32  1 1 SE     -
52480
22.22.22.22     58.58.58.2      5001     2205     DTR    TR-PE1-PE2-MP-01-Detour         UP    00:34:04  1 2 SE     -
25600
Total 2 displayed, Up 2, Down 0.

Transit RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     61.61.61.2      5001     2205     PRI    TR-PE1-PE2-MP-01-Detour         UP    00:33:19  1 2 SE     25602
25600
Total 1 displayed, Up 1, Down 0.

Egress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
26.26.26.26     22.22.22.22     5001     2205     PRI    TR-PE2-PE1-MP-01-Primary        UP    02:12:27  1 1 SE     25601    -
26.26.26.26     62.62.62.2      5001     2205     PRI    TR-PE2-PE1-MP-01-Detour         UP    02:09:08  1 1 SE     25600    -
Total 2 displayed, Up 2, Down 0.
```

Below is the validation output of RSVP ping and trace from PE1 to PE2:

```
#ping mpls rsvp egress 22.22.22.22 detail
Sending 5 MPLS Echos to 22.22.22.22, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
```

```
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

! seq_num = 1 56.56.56.3 0.91 ms
! seq_num = 2 56.56.56.3 0.54 ms
! seq_num = 3 56.56.56.3 0.48 ms
! seq_num = 4 56.56.56.3 0.47 ms
! seq_num = 5 56.56.56.3 0.50 ms

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.47/0.69/0.91
PE1#
#trace mpls rsvp egress 22.22.22.22 detail
Tracing MPLS Label Switched Path to 22.22.22.22, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

  0 61.61.61.3 [Labels:  52480]
R 1 61.61.61.2 [Labels:  25600] 0.71 ms
R 2 23.23.23.3 [Labels:  25600] 0.83 ms
R 3 41.41.41.3 [Labels:  25600] 0.88 ms
! 4 56.56.56.3 0.69 ms
```

Below are the outputs from transit nodes R1, R2 and R3 for primary LSP configured:

R1

```
#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                        State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     61.61.61.2      5001     2205     DTR    TR-PE1-PE2-MP-01-Detour        UP    00:38:43  1 2 SE     -
25602
26.26.26.26     23.23.23.2      5001     2205     DTR    TR-PE2-PE1-MP-01-Detour        UP    00:38:44  1 1 SE     -
25603
Total 2 displayed, Up 2, Down 0.

Transit RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                        State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     26.26.26.26     5001     2205     PRI    TR-PE1-PE2-MP-01-Primary       UP    02:17:55  1 1 SE     52480
25600
22.22.22.22     23.23.23.3      5001     2205     PRI    TR-PE1-PE2-MP-01-Detour        UP    00:37:58  1 2 SE     52482
25602
26.26.26.26     22.22.22.22     5001     2205     PRI    TR-PE2-PE1-MP-01-Primary       UP    02:17:50  1 1 SE     52481
25601
Total 3 displayed, Up 3, Down 0.
```

R2

```
#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
```

```
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     23.23.23.3      5001     2205     DTR    TR-PE1-PE2-MP-01-Detour         UP    00:38:07  1 2 SE     -
52482
26.26.26.26     41.41.41.2      5001     2205     DTR    TR-PE2-PE1-MP-01-Detour         UP    00:39:00  1 2 SE     -
25602
Total 2 displayed, Up 2, Down 0.

Transit RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     26.26.26.26     5001     2205     PRI    TR-PE1-PE2-MP-01-Primary        UP    02:18:05  1 1 SE     25600
25600
22.22.22.22     41.41.41.3      5001     2205     PRI    TR-PE1-PE2-MP-01-Detour         UP    00:37:28  1 2 SE     25602
52482
26.26.26.26     22.22.22.22     5001     2205     PRI    TR-PE2-PE1-MP-01-Primary        UP    02:18:00  1 1 SE     25601
52481
26.26.26.26     23.23.23.2      5001     2205     PRI    TR-PE2-PE1-MP-01-Detour         UP    00:38:53  1 2 SE     25603
25602
Total 4 displayed, Up 4, Down 0.
```

R3

```
#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     41.41.41.3      5001     2205     DTR    TR-PE1-PE2-MP-01-Detour         UP    00:37:31  1 1 SE     -
25602
26.26.26.26     56.56.56.2      5001     2205     DTR    TR-PE2-PE1-MP-01-Detour         UP    00:39:23  1 2 SE     -
25602
Total 2 displayed, Up 2, Down 0.

Transit RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     26.26.26.26     5001     2205     PRI    TR-PE1-PE2-MP-01-Primary        UP    02:18:08  1 1 SE     25600
25600
26.26.26.26     22.22.22.22     5001     2205     PRI    TR-PE2-PE1-MP-01-Primary        UP    02:18:02  1 1 SE     25601
25601
26.26.26.26     41.41.41.2      5001     2205     PRI    TR-PE2-PE1-MP-01-Detour         UP    00:39:03  1 2 SE     25602
25602
Total 3 displayed, Up 3, Down 0.
```

Below are the outputs from transit nodes R4 and R5 for Detour LSPs formation:

From R4

```
#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Transit RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     58.58.58.2      5001     2205     PRI    TR-PE1-PE2-MP-01-Detour         UP    02:14:52  1 1 SE     25600
25601
26.26.26.26     62.62.62.2      5001     2205     PRI    TR-PE2-PE1-MP-01-Detour         UP    00:39:49  1 1 SE     25601
25601
Total 2 displayed, Up 2, Down 0.
```

From R5

```
#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Transit RSVP:
```

```
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     58.58.58.2      5001     2205     PRI    TR-PE1-PE2-MP-01-Detour         UP    00:39:45  1 1 SE     25600
25600
26.26.26.26     62.62.62.2      5001     2205     PRI    TR-PE2-PE1-MP-01-Detour         UP    02:14:48  1 1 SE     25601
25600
Total 2 displayed, Up 2, Down 0.
```

Now, shutting down one of the interfaces on Primary LSP path and check RSVP tunnel outputs on PE1 and PE2

Shutdown interface xe47 connected between R1 and R2:

`#configure terminal`	Enter Configure mode.
`(config)#interface xe47`	Enter interface mode.
`(config-router)#shutdown`	Administratively bring the interface down.
`(config-router)#exit`	Exit router RSVP mode

Below is the validation output of RSVP LSPs from PE1 to PE2 after admin shutting one of the interfaces on primary LSP path:

```
#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     26.26.26.26     5001     2205     PRI    TR-PE1-PE2-MP-01-Primary        UP*   02:32:40  1 1 SE     -
52480
22.22.22.22     26.26.26.26     5001     2201     PRI    TR-PE1-PE2-MP-01-Primary        DN    N/A       0 0 SE     -         -
22.22.22.22     58.58.58.2      5001     2205     DTR    TR-PE1-PE2-MP-01-Detour         UP    00:54:12  1 2 SE     -
25600
Total 3 displayed, Up 2, Down 1.

Transit RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     61.61.61.2      5001     2205     PRI    TR-PE1-PE2-MP-01-Detour         UP    00:53:27  1 2 SE     25602
25600
Total 1 displayed, Up 1, Down 0.
```

Below is the validation output of RSVP ping and trace from PE1 to PE2 after shutting one of the interfaces on primary LSP path:

```
Egress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
26.26.26.26     62.62.62.2      5001     2205     PRI    TR-PE2-PE1-MP-01-Detour         UP    02:29:16  1 1 SE     25600     -
Total 1 displayed, Up 1, Down 0.

#ping mpls rsvp egress 22.22.22.22 detail
Sending 5 MPLS Echos to 22.22.22.22, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

! seq_num = 1 62.62.62.2 0.69 ms
! seq_num = 2 62.62.62.2 0.54 ms
! seq_num = 3 62.62.62.2 0.56 ms
```

```
! seq_num = 4 62.62.62.2 0.49 ms
! seq_num = 5 62.62.62.2 0.51 ms

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.49/0.59/0.69
#trace mpls rsvp egress 22.22.22.22 detail
Tracing MPLS Label Switched Path to 22.22.22.22, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

  0 61.61.61.3 [Labels:  52480]
R 1 61.61.61.2 [Labels:  25602] 0.72 ms
R 2 61.61.61.3 [Labels:  25600] 0.67 ms
R 3 58.58.58.3 [Labels:  25600] 0.80 ms
R 4 54.54.54.3 [Labels:  25601] 0.80 ms
! 5 62.62.62.2 0.50 ms
```

Below is the validation output of RSVP LSPs from PE2 to PE1 after admin shutting one of the interfaces on primary LSP path:

```
#show rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
26.26.26.26     22.22.22.22     5001     2205     PRI    TR-PE2-PE1-MP-01-Primary        UP*   02:36:19  1 1 SE     -
25601
26.26.26.26     22.22.22.22     5001     2201     PRI    TR-PE2-PE1-MP-01-Primary        DN    N/A       0 0 SE     -        -
26.26.26.26     62.62.62.2      5001     2205     DTR    TR-PE2-PE1-MP-01-Detour         UP    00:57:57  1 2 SE     -
25601
Total 3 displayed, Up 2, Down 1.

Transit RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
26.26.26.26     56.56.56.2      5001     2205     PRI    TR-PE2-PE1-MP-01-Detour         UP    00:57:40  1 2 SE     25602
25601
Total 1 displayed, Up 1, Down 0.

Egress RSVP:
To              From            Tun-ID   LSP-ID   Type   LSPName                         State Uptime    Rt  Style  Labelin
Labelout
22.22.22.22     58.58.58.2      5001     2205     PRI    TR-PE1-PE2-MP-01-Detour         UP    02:33:00  1 1 SE     25601    -
Total 1 displayed, Up 1, Down 0.
```

Below is the validation output of RSVP ping and trace from PE2 to PE1 after shutting one of the interfaces on primary LSP path:

```
#ping mpls rsvp egress 26.26.26.26 detail
Sending 5 MPLS Echos to 26.26.26.26, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort
```

```
! seq_num = 1 58.58.58.2 0.80 ms
! seq_num = 2 58.58.58.2 0.59 ms
! seq_num = 3 58.58.58.2 0.47 ms
! seq_num = 4 58.58.58.2 0.49 ms
! seq_num = 5 58.58.58.2 0.54 ms

Success Rate is 100.00 percent (5/5)
round-trip min/avg/max = 0.47/0.63/0.80
#trace mpls rsvp egress 26.26.26.26 detail
Tracing MPLS Label Switched Path to 26.26.26.26, timeout is 5 seconds

Codes:
'!' - Success, 'Q' - request not sent, '.' - timeout,
'x' - Retcode 0, 'M' - Malformed Request, 'm' - Errored TLV,
'N' - LBL Mapping Err, 'D' - DS Mismatch,
'U' - Unknown Interface, 'R' - Transit (LBL Switched),
'B' - IP Forwarded, 'F' No FEC Found, 'f' - FEC Mismatch,
'P' - Protocol Error, 'X' - Unknown code,
'Z' - Reverse FEC Validation Failed

 Type 'Ctrl+C' to abort

  0 56.56.56.3 [Labels:  25601]
R 1 56.56.56.2 [Labels:  25601] 1.01 ms
R 2 41.41.41.2 [Labels:  25602] 0.95 ms
R 3 41.41.41.3 [Labels:  25602] 0.62 ms
R 4 56.56.56.3 [Labels:  25601] 0.79 ms
R 5 62.62.62.3 [Labels:  25601] 0.67 ms
R 6 54.54.54.4 [Labels:  25600] 0.57 ms
! 7 58.58.58.2 0.50 ms
```

Implementation Examples

To implement detour based protection in a ring topology, use the command detour-allow-primary-upstream-path that allows the detour formation to consider the upstream path of protected LSP. This is only applicable in ring topology.

New CLI Commands

detour-allow-primary-upstream-path

Use this command to ensure detour formation to consider the upstream path of protected LSPs. This is a deviation to RFC 4090 section 6.2 recommendation (https://datatracker.ietf.org/doc/html/rfc4090). This command is intended to be used in special cases where detour protection is required on ring topology if no alternate path is available.

Use the no parameter with this command to bypass the upstream path to the protected LSP when choosing a detour path.

Note: This command is intended to be used in ring topology if detour support is required at the cost of resource and link bandwidth. This command is not recommended to be configured otherwise.

Command Syntax

```
detour-allow-primary-upstream-path
no detour-allow-primary-upstream-path
```

Parameters

None

Default

By default, detour formation excludes the protected LSP upstream path as per RFC 4090 section 6.2 recommendations.

Command Mode

Router mode

Applicability

This command was introduced in OcNOS version 6.4.1.

Examples

```
#configure terminal
(config)#router rsvp
(config-router)#detour-allow-primary-upstream-path
(config-router)#commit
(config-router)#no detour-allow-primary-upstream-path
(config-router)#commit
```

Abbreviations

The following are some key abbreviations and their meanings relevant to this document:

Acronym	Description
FRR	Fast Reroute
LSP	Label Switched Path
OSPF	Open Shortest Path First
PLR	Point of Local Repair

Glossary

The following provides definitions for key terms used throughout this document:

Detour formation in the ring topology	The detour formation in the ring topology is a mechanism to reroute the data traffic over the backup path when a failure or congestion occurs in the primary Label Switched Path (LSP).
PLR node	A PLR node is a network device that reacts and takes action when a link fails.
Primary LSP	The primary LSP is the default path of the forwarding data packets from the source device to the destination device.
Protected LSP	A protected LSP is a primary LSP with a backup path in an MPLS network. When there is an issue or a failure in the primary LSP, the traffic is rerouted through the backup path, protecting the primary LSP.

RSVP Tunnel	RSVP tunnels are logical paths through which data traffic traverses in an IP network.
Upstream path of the protected LSP	The upstream path of the protected LSP is the section of the network that precedes the PLR node in the network.

Commit Rollback

Overview

The Commit Rollback capability in Common Management Layer Commands (CMLSH) is designed to execute a rollback operation for a set of configurations that were previously committed, with each commit operation identified by a unique commit ID. The Commit ID is numeric value and is generated by the CMLSH Commit, Confirmed Commit and Commit Rollback.

This Commit Rollback application is used for rolling back the commits that are performed after the specified commit ID whether they were executed through either Commit or Confirmed Commit operations.

Here, you find the description for Commit and Confirmed Commit:

- **Commit operation**: Involves committing the candidate configuration to the running configuration.
- **Confirmed Commit operation**: Provides more options to the commit operation with timeout parameter, user could provide timeout for the commit (default is 300 seconds).

During this timeout interval, users can either confirm the commit or cancel it, and if no confirmation or cancellation is provided before the timer expires, commit will be automatically rolled back after timeout. For an example, see the Example section of commit-rollback CLI.

Feature Characteristics

The Confirmed-Commit operation temporarily applies the configuration for the duration specified in seconds. If the user does not confirm the configuration within this timeframe, an automatic rollback will be initiated once the timer expires. For committing the configurations with timings, see commit.

Once the configurations are confirmed, users can use the commit rollback operation to revert the configuration, whether it is for a commit operation or a confirmed commit operation.

Benefits

With the integration of CMLSH Commit Rollback with Standard or Confirmed Commit, users can initiate a rollback operation for any specific commit, utilizing the associated commit ID to revert the configurations to their previous state. In this way, reverting to an earlier state, functional configuration is possible in case the new configuration is compromised or if the configuration makes the device unstable.

Prerequisites

Before configuring this operation, enable cml commit-history to ensure the commit records are stored in the commit history list. By default, `cml commit-history` is enabled. For enabling or disabling it, see cml commit-history (enable | disable).

Commands for Commit Rollback

Note: For the commands, refer to the Common Management Layer Commands section in the *System Management Command Reference guide*.

Abbreviations

List of key terms used in this document is:

Term	Description
CMLSH	Common Management Layer Commands

Index