ENTERPRISE LAYER 3 MULTICAST IMPLEMENTATION GUIDE

Using the EX Series Ethernet Switches and MX Series 3D Universal Edge Routers

Table of Contents

Table of Figures

Introduction

Enterprise customers are increasingly deploying IP multicast forwarding in their networks to deliver applications such as video conferencing, distance learning, and distribution of software, stock quotes, and news. They use Protocol Independent Multicast (PIM) for multicast signaling and the Internet Group Management Protocol (IGMP) to manage host multicast group membership.

The objective of this document is to provide a Layer 3 multicast implementation guide for enterprise networks using the Juniper Networks® EX Series Ethernet Switches and Juniper Networks MX Series 3D Universal Edge Routers. This document targets deployments where the EX Series is used in the access layer, while core and aggregation are collapsed into one layer using the MX Series. It first provides general design considerations and compares different options for PIM and IGMP deployments. It then provides implementation guidelines with configuration procedures. Finally, an implementation example is included along with the network topology, detailed configurations, as well as verification and troubleshooting procedures.

Scope

This guide focuses on implementing PIM and IGMP in a Layer 3 enterprise environment using the EX Series and MX Series. The Juniper Networks EX4200 Ethernet Switch and Juniper Networks EX3200 Ethernet Switch are used as Layer 3 and Layer 2 devices with virtual LAN (VLAN) interfaces, PIM, and IGMP enabled. IGMP snooping is also enabled. OSPF is used but is not the main focus of this implementation.

This document does not discuss Distance Vector Multicast Routing Protocol (DVMRP), Multicast Source Discovery Protocol (MSDP), Session Announcement Protocol (SAP), or multicast VPN. It also does not cover IPv6.

Target Audience

This document is intended for network design, operation engineers, or other technical audiences seeking IP multicast implementation guidelines for enterprise deployments using the EX Series Ethernet Switches and MX Series 3D Universal Edge Routers.

Design Considerations

Choosing the PIM Mode

- PIM dense mode (PIM DM): Implements a flood-and-prune mechanism to build a source-based distribution tree. A router receives the multicast data on the interface closest to the source and floods the traffic to all other interfaces. Routers with no receivers must prune back unnecessary branches.
- PIM sparse mode (PIM SM): Uses reverse path forwarding (RPF) to create a path from a data source to the receiver requesting the data. When a receiver issues an explicit join request, an RPF check is triggered. A single router called a rendezvous point (RP) is initially selected in each domain to be the connection point between sources and interested clients. Traffic flows are then rooted at the RP along the rendezvous-point tree. The rendezvous-point tree is later replaced by optimized shortest-path tree.
- PIM source-specific multicast (PIM SSM): Uses a subset of PIM sparse mode and IGMP version 3 to allow a client to receive multicast traffic directly from the source. PIM source-specific multicast builds the shortest-path tree between the receiver and the source, but without the help of a RP.

Table 1: PIM Modes Comparison

MODE	PROS	CONS
PIM-DM	• No requirement for RPs • Guarantees shortest path from source to the receiver	• Considerable overhead • Does not scale well
PIM-SM	• Better scalability than PIM dense mode	• Requirement for RPs • May require special hardware to encapsulate/de-encapsulate register messages (depending on the platforms)
PIM-SSM	• No requirement for RPs • More secure and prevents data plane attacks where malicious hosts flood unwanted traffic on a group • No shared tree behavior	• Requires IGMP version 3

Selecting and Assigning RPs

It is important to select reliable devices to act as RPs and carefully choose their location in order to improve the performance and fault tolerance of an enterprise network using PIM sparse mode. The RPs can be assigned statically or dynamically. To provide RP redundancy, it is recommended to choose a dynamic RP discovery method. The following table lists the options for RP assignment and discovery and compares their pros and cons.

Table 2: RP Discovery Methods Comparison

MODE	PROS	CONS
Static RP	• Simplicity • Operates with PIM version 1 and 2	• No fault tolerance • Single point of failure
Auto-RP	• Dynamic method • Redundancy • Failover mechanism • Operates with both PIM version 1 and 2	• Non-standard (Cisco proprietary) • Requires use of dense-mode groups to advertise control traffic (224.0.1.39 for announce messages and 224.0.1.40 for discovery messages) • Slower failover as opposed to anycast RP because of time involved in noticing the failure and electing a new RP for that group • One RP is operational at a time
Bootstrap	• Dynamic method • Redundancy • Failover mechanism • Standardized method (as opposed to auto-RP) • Does not require dense-mode groups for control traffic • Multiple routers can be candidates BSR or candidate RPs	• Slower failover as opposed to anycast RP because of time involved in noticing the failure and electing a new RP for that group • One RP is operational at a time
Anycast RP	• Virtual RP for the entire domain • Multiple physical routers share knowledge about multicast sources • Fastest convergence around failures • Best load-balancing and redundancy	• Requires MSDP (unless using anycast with PIM only for a single domain)

Notes:

- Juniper Networks Junos® operating system prefers RPs learned through bootstrap over auto-RP. Both dynamic options are preferred over a statically configured RP.

Selecting the IGMP Version

The table below provides a comparison of IGMP versions supported by Junos OS:

Table 3: IGMP Versions Comparison

VERSION	RFC	PROS	CONS
Version 1	**1112**	· Periodic host membership query messages sent by all routers to all-hosts group address (224.0.0.1)	· High latency: since there is no leave-group mechanism, multicast traffic can continue to be forwarded for several minutes after the last host leaves the group
Version 2	**2236**	· Adds querier election process: the lowest IP address on the LAN is selected · Defines group-specific query and explicit leave-group messages	· Improved latency compared to version 1
Version 3	**3376**	· Enhances version 2 support of leave-group messages by introducing group-and-source specific report messages · Accommodates source-specific multicast (SSM)	· Efficiency gain compared to version 2 · Can also support a sparse-mode topology without a rendezvous point · Hosts must have a priori knowledge of the specific sources active for a given group

Note: Junos OS defaults to IGMP version 2.

Reducing Multicast Flooding with IGMP Snooping

To avoid flooding of all multicast traffic on a VLAN, IGMP snooping should be enabled on access switches so they intercept IGMP packets. The switch then uses the content of the packets to build a multicast cache table that is a database of multicast groups and their corresponding member ports. This table regulates multicast traffic on the VLAN.

Note: Factory-deafult configuration on EX Series switches enables IGMP snooping on all vlans using “set protocols igmp-snooping vlan all” command.

Reducing IGMP Membership Packets with IGMP Proxy

In deployments where an aggregation switch is an IGMP query and the access switch has many hosts, sending IGMP queries to all hosts results into a storm of IGMP membership packets. IGMP proxy allows the access switch to prevent this by generating IGMP reports on behalf of the hosts. The switch also generates periodic IGMP query to hosts to maintain its list of group/port membership. When a host sends an IGMP join message, the switch suppresses this report (and updates its timer) if this is not the first join for the group on that VLAN. Similarly, leave messages received from a host are suppressed unless it is the last leave for the group on the VLAN.

Note: IGMP snooping proxy is not supported on EX Series Ethernet Switches as of release 9.3.

Hardware Considerations

In PIM sparse mode, the source designated router takes the initial multicast packets and encapsulates them in PIM register messages. It then unicasts them to the PIM sparse-mode RP router, where the PIM register message is de-encapsulated. Therefore, RP routers and designated routers connected to a source require special hardware (Tunnel Services PIC) to encapsulate and de-encapsulate PIM register messages.

If the source designated router is the RP, then there is no requirement for PIM register messages and consequently no requirement for a Tunnel Services PIC to be present.

On EX Series Ethernet Switches, Tunnel Services PICs are not required to encapsulate and de-encapsulate register messages since this is done by the RE CPU.

On MX Series 3D Universal Edge Routers, there is no Tunnel Services PIC but Tunnel Services PIC functionality can be enabled by configuring the chassis as shown in the example below:

Enabling Tunnel PIC functionality on MX Services 3D Universal Edge Routers:

```
chassis {
        fpc 2 {
                pic 0 {
                        tunnel-services {
                                bandwidth 1g;
                        }
                }
        }
}
```

Scaling Numbers on EX3200 and EX4200 Platforms

- Max L3 multicast routing table entries: 2K
- Max L2 multicast entries (IGMP snooping): 8K
- 800 joins/second tested by system test

EX Series Multicast Feature Support Summary

The table below lists the IGMP and PIM features supported on EX3200 and EX4200 and shows the software releases they were introduced:

Table 4: Multicast Features Support on EX3200 and EX4200

	FEATURE	EX3200	EX4200
IGMP	V1/V2	9.0	9.0
	V3	9.3	9.3
	V1/V2 snooping	9.1	9.1
PIM	PIM-SM	9.0	9.0
	PIM-DM	9.2	9.2
	PIM-SSM	9.3	9.3

Implementation

Implementation Guidelines

IGMP Configuration

Enabling/Disabling IGMP

PIM is required on upstream IGMP interfaces to distribute IGMP group memberships into the multicast routing domain. By default, IGMP is automatically enabled on all interfaces where PIM is configured. It is also possible to enable/disable IGMP explicitly on an interface with the following statements:

Enabling IGMP explicitly on an interface:

```
protocols{
        igmp {
                interface interface-name;
        }
}
```

Disabling IGMP explicitly on an interface:

```
protocols{
        igmp {
           interface interface-name;
           disable;
        }
}
```

Note: To enable/disable IGMP on all interfaces at once, use the same commands above, replacing the interface name with the keyword “all”.

IGMP Version

The default IGMP version applied by Junos OS is version 2. The version used can be changed with the following statement. Note that this can be done also at the interface level.

Changing the IGMP version:

```
protocols{
   igmp {
       version version-number;
   }
}
```

Note: If two routers run different versions of IGMP, they negotiate the lowest common version of IGMP that is supported by hosts on their subnet.

IGMP Timers

The querier interval is the time between periodic host-query messages sent to the all-systems IP address of 224.0.0.1. It is set by default to 125 seconds but can be modified to a value from one to 1024 seconds using the following statement:

Changing the IGMP query-interval:

```
protocols{
   igmp {
       query-interval seconds;
   }
}
```

The query-response-interval indicates the maximum duration between when the querier router sends a host-query message and when it receives a response from a host. To adjust the burst peaks of IGMP messages on a subnet, this timer can be modified from its default value of 10 seconds to a value between one and 1024 seconds as shown below:

Changing the IGMP query-response-interval:

```
protocols{
   igmp {
        query-response-interval seconds;
   }
}
```

IGMP Immediate-Leave

On IGMP version 2 interfaces that have only one host connected, the router can be configured to remove a host from the multicast group immediately after receiving a leave group message. An example is shown below:

```
Applying immediate-leave to an interface:
    protocols{
          igmp {
                  interface interface-name {
                            immediate-leave;
                  }
           }
     }
```

Note: This should not be applied to interfaces attached to multiple hosts since the router would remove all hosts until they send join requests in response to the router's next general group membership query.

IGMP Promiscuous Mode

You can allow a router to accept IGMP messages from indirectly connected subnets (for example from sources that do not match the IP subnet of the interface). An interface is set to promiscuous mode using the following statements:

```
Applying the promiscuous-mode to an interface:
    protocols{
        igmp {
        interface interface-name {
                              promiscuous-mode;
        }
        }
    }
```

Static IGMP

You can configure static IGMP to help test and verify multicast forwarding in the absence of receivers. The static join can take the form of a (,G) or (S,G) entry, based on the inclusion of a source address. If a source address is specified, the IGMP version must be set to IGMPv3.

```
Configuring static IGMP:
        protocols{
           igmp {
                   interface interface-name {
                          static {
                                  group group-address {
                                      source address;
                                   }
                              }
                        }
                }
           }
```

IGMP Snooping

You can configure IGMP snooping:

- On all VLans, which is the factory-default configuration for EX switches.
- On individual interfaces in the VLAN
- Statically for a group on an individual interface

Enabling IGMP snooping in a VLAN:

```
     protocols{
          igmp-snooping {
               vlan vlan-id;
               }
          }
        }
```

Enabling IGMP snooping on a specific interface:

```
        protocols{
           igmp-snooping {
                   vlan vlan-id
                          interface interface-name;
                   }
        }
    }
```

Statically configuring IGMP snooping for a specific group on an interface:

```
    protocols{
        igmp-snooping {
           vlan vlan-id {
                  interface interface-name {
                         static {
                                group group-address;
                         }
                     }
                  }
              }
        }
```

It is possible to configure the switch to immediately remove a multicast group membership from an interface when it receives a leave message from that interface and suppresses the sending of any group-specific queries for the multicast group (IGMPv2 only) as shown below:

Setting IGMP snooping on a VLAN to immediate leave:

```
    protocols{
        igmp-snooping {
                vlan vlan-id {
                        immediate-leave;
                }
        }
    }
```

Other configuration options include changing the query-last-member, query-response, or query-interval on a VLAN as shown in the example below:

Changing the IGMP snooping query interval on a VLAN:

```
    protocols{
        igmp-snooping {
                vlan vlan-id {
                       query-interval seconds;
                       }
                }
```

```
        }
    }
```

PIM Configuration

Static RP

If using static-RP, each router in the PIM domain must be configured with the location of the RP. The address of the RP is specified, as well as the designated RP group. If a group is not specified, the RP is used for all possible group addresses (for example, 224.0.0.0/4).

Configuring Static-RP:

```
      protocols{
            pim {
                rp {
                   static {
                          address address
                          group-ranges {
                              destination-mask;
                          }
                    }
                }
            }
      }
```

Auto-RP

Auto-RP is configured by completing the following steps on all routers in the network:

Step 1: Enabling PIM sparse-dense mode on all interfaces

Step 2: Enabling PIM dense mode for the announce (224.0.1.39) and discovery (224.0.1.40) groups

Step 3: Enabling auto-RP with one of the following options:

- Discovery: allows the router to listen for mapping messages
- Announce: allows the router to listen for mapping messages and advertise that it can be an RP
- Mapping: in addition to the discovery and announce capability, allows the router to elect the RP and send group-to-RP mapping and discovery messages

Auto-RP configuration on all routers:

```
    protocols{
        pim {
           dense-groups{
              224.0.1.39/32;
              224.0.1.40/32;
           }
           rp {
              auto-rp {announce | discovery | mapping};
           }
           interface all {
                mode sparse-dense;
           }
           interface me0.0 {
               disable;
           }
      }
   }
```

RP local address configuration:

```
protocols{
    pim {
        rp {
           local {
               address address;
           }
        }
     }
 }
```

Bootstrap RP

A bootstrap router (BSR) is configured by completing the following steps:

Step 1: Configuring one or more routers as candidate BSRs by setting the bootstrap priority at the PIM RP level.

Step 2: Selecting one or more routers as candidate RPs by configuring the local address of the RP at the PIM RP level.

Configuring C-BSR:

```
protocols{
    pim {
       rp {
          bootstrap-priority priority;
       }
    }
 }
```

Configuring C-RP:

```
protocols{
    pim {
        rp {
           local {
              address address;
           }
        }
    }
 }
```

Note: The router with the highest priority becomes the BSR for the PIM domain. Should the priority of two routers be equal, the one with the highest IP address is selected. Once the bootstrap router is selected in the PIM domain, each router with a local RP configuration in the domain advertises its capabilities in a candidate RP-Adv message that is unicast to the BSR.

Anycast RP with PIM Only

Anycast RP can be enabled without MSDP using the anycast-PIM configuration. The RP routers that share the same IP address are configured using the RP-set statement. Below are the steps and configuration statements:

Step 1: On all RP routers, configure the shared anycast address on the loopback interface to allow it to be advertised by the IGP. The "primary" keyword is included with the unicast address configuration to ensure that the automatic selection of a router ID returns the unicast value and not the shared anycast value.

Step 2: On all RP routers, specify the shared anycast RP address using the address statement under "pim rp".

Step 3: On all RP routers, specify the unicast IP address of the local RP using the local-address statement under "pim rp local anycast-pim."

Step 4: On all RP routers, specify the unicast IP addresses of other RP routers using the rp-set statement under "pim rp local anycast-pim."

Step 5: On non-RP routers, configure a static RP using the shared anycast address.

Step 6: For all interfaces, use the mode statement to set the mode to sparse. When configuring all interfaces, exclude

the management interface by adding the disable statement for that interface.

Configuring RP routers for anycast RP with PIM only:

```
   interface{
      lo0 {
            unit 0 {
                family inet {
                       address Unicast-address/mask {
                           primary;
                       }
                       address shared-anycast-address/mask;
                }
           }
     }
}
protocols {
    pim {
           rp {
                 local {
                        family inet {
                               address shared-anycast-address;
                               anycast-pim {
                                      rp-set {
                                             address Unicast-address;
                               }
                               local-address Unicast-address;
                         }
                    }
               }
         }
         interface all {
                mode sparse;
         }
         interface fxp0.0 {
                 disable;
         }
   }
}
```

Configuring non-RP routers for anycast RP with PIM only:

```
    protocols {
        pim {
              rp {
                       static {
                              address shared-anycast-address;
                        }
               }
        }
     }
```

Notes

- If the local-address statement is omitted, the primary loopback address is used.
- The maximum number of routers that can be configured as RPs is 15.

Implementation Example

Topology

Figure 1: Network topology

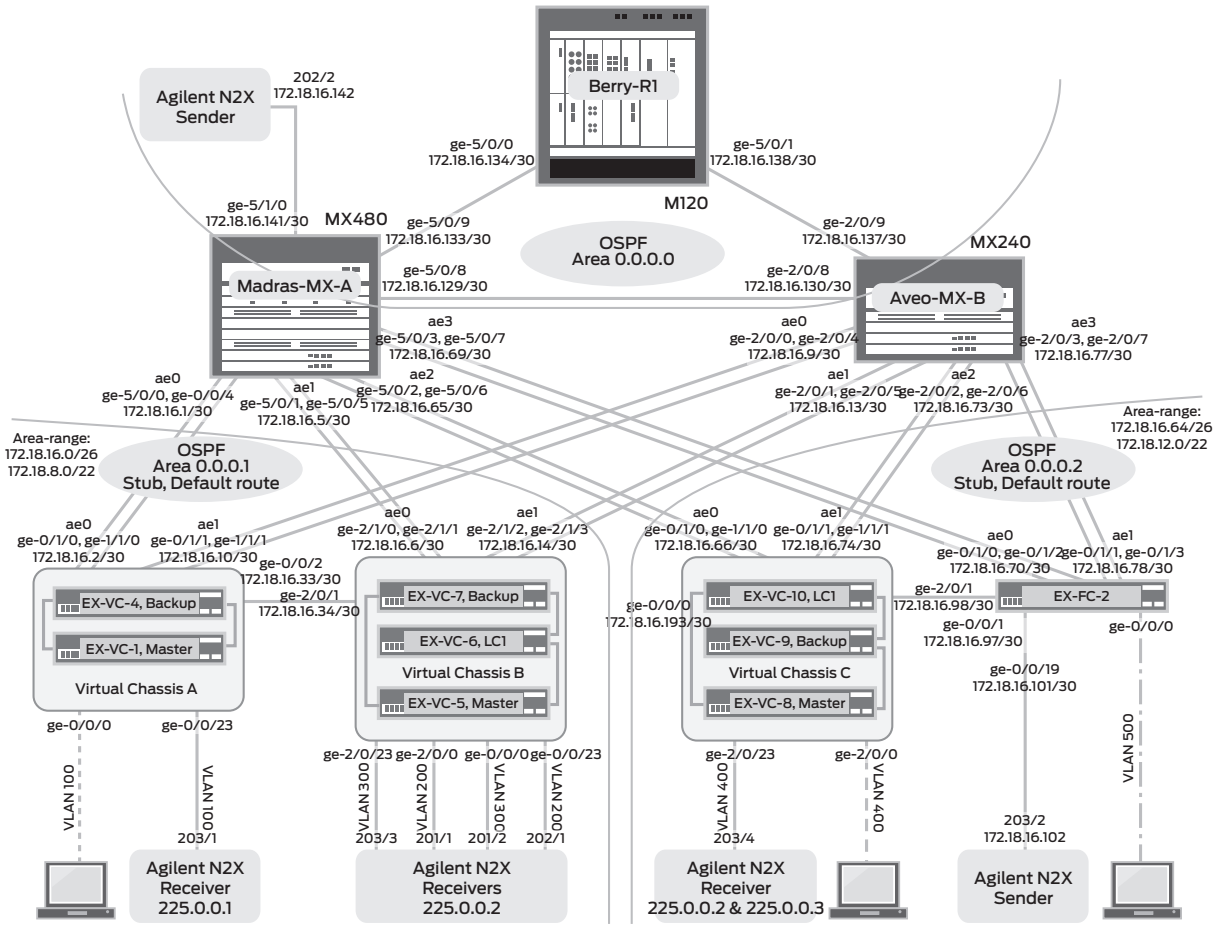

RVIs:
Virtual Chassis A: vlan.100: 172.18.9.1/24
Virtual Chassis B: vlan.200: 172.18.10.1/24
vlan.200: 172.18.10.1/24
Virtual Chassis C: vlan.400: 172.18.13.1/24
Fixed Configuration Chassis: vlan.500: 172.18.14.1/24

Loopbacks:
MX-A: 172.18.19.1 MX-B: 172.18.19.2
R1: 172.18.19.3 EX-VC-1: 172.18.8.1
EX-VC-5: 172.18.8.2 EX-VC-8: 172.18.12.1
EX-FC-2: 172.18.12.2

Static Routes:
MX-A and MX-B: 200.0.1.0/24 and 200.0.2.0/24

OSPF Authentication:
Area): MD5 key=juniper), key-id=10
Area1Z: Text password = juniper1
Area2: Text password = juniper2

N2X multicast traffic streams:

	Source port -> Destination
StreamGroup 9:	port 202/2 -> 225.0.0.1
StreamGroup 10:	port 202/2 -> 225.0.0.2
StreamGroup 11:	port 202/2 -> 225.0.0.3
StreamGroup 12:	port 203/2 -> 225.0.0.1
StreamGroup 13:	port 203/2 -> 225.0.0.2
StreamGroup 14:	port 202/2 -> 225.0.0.3

Hardware Used For the Implementation

Table 5: Implementation Example—Hardware Requirements

EQUIPMENT	COMPONENTS
4 x EX4200 Ethernet Switch	• 4 x 4-port uplink 1-Gigabit Ethernet module (EX-UM-4SFP) • 16 SFPs • 8 x VCP cables
1 x MX240 3D Universal Edge Router **1 x MX480**	• 2 x 40-port 1-Gigabit Ethernet L2/L3 DPCs (DPCE-R-40GE-SFP or DPCE-R-Q-40GE-SFP) • 20 SFPs
1 x M120 Multiservice Edge Router	• 1 x FPC Type 3 • 1 x 10-port 1-Gigabit Ethernet PIC • 2 SFPs
Agilent N2X tester	• 10 x 1-Gigabit Ethernet ports (9x RJ45 and 1 SFP type) • 1 SFP

Software Used For the Implementation

In this topology, PIM-SM is used and MX-A and MX-B are configured as RPs using anycast without MSDP. IGMP version 2 along with IGMP snooping are enabled on all EX Series Ethernet Switches connected to multicast receivers.

Table 6: Implementation Example—Software Requirements

EQUIPMENT	MULTICAST FEATURES	SOFTWARE
EX Series Ethernet Switches	• IGMPv2 • IGMP snooping • PIM-SM • Anycast RP	Junos OS 9.3
MX Series 3D Universal Edge Routers and M Series Multiservice Edge Routers	• PIM-SM • Anycast RP	

Detailed Configurations

EX-VC-1 Configuration

```
#…truncated…
chassis {
    redundancy {
        graceful-switchover;
    }
    aggregated-devices {
        ethernet {
            device-count 2;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members HR;
                }
```

```
        }
    }
}
ge-0/0/2 {
    unit 0 {
        family inet {
            address 172.18.16.33/30;
        }
    }
}
ge-0/0/23 {
    unit 0 {
        family ethernet-switching {
            port-mode access;
            vlan {
                members HR;
            }
        }
    }
}
ge-0/1/0 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae0;
    }
}
ge-0/1/1 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae1;
    }
}
ge-1/1/0 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae0;
    }
}
ge-1/1/1 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae1;
    }
}
ae0 {
    unit 0 {
        family inet {
            address 172.18.16.2/30;
        }
    }
}
ae1 {
    unit 0 {
        family inet {
            address 172.18.16.10/30;
        }
    }
}
```

```
    lo0 {
        unit 0 {
            family inet {
                address 172.18.8.1/32;
            }
        }
    }
    vlan {
        unit 100 {
            family inet {
                address 172.18.9.1/24;
            }
        }
    }
    vme {
        unit 0 {
            family inet {
                address 172.19.59.190/24;
            }
        }
    }
}
routing-options {
    graceful-restart;
    router-id 172.18.8.1;
}
protocols {
    igmp;  # Explicitly enables IGMP
    ospf {
        area 0.0.0.1 {
            stub default-metric 10;
            interface lo0.0 {
                passive;
            }
            interface vme.0 {
                disable;
            }
            interface vlan.100 {
                passive;
            }
            interface ae0.0 {
                authentication {
                    simple-password "$9$KRdWX-YgJHqfVwqfTzCAvWLxVw"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface ae1.0 {
                authentication {
                    simple-password "$9$pUpQOIcKMXbs4yls4aZkquO1Ryl"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface ge-0/0/2.0 {
                authentication {
                    simple-password "$9$JgUi.QF/0BEP5BEcyW8ZUjHP5"; ## SECRET-DATA
                }
                bfd-liveness-detection {
```

```
                    minimum-interval 300;
                }
            }
        }
    }
    pim {
        rp {   # Configuration of anycast-PIM on non-RP devices
            static {
                address 172.18.19.254;  # Uses the shared anycast address as the RP address
            }
        }
        interface all {
            mode sparse;  # Sets all interfaces to PIM sparse mode
        }
        interface vme.0 {
            disable;  # Explicitly disables PIM on the management interface
        }
    }
    igmp-snooping {
        vlan HR;  # Enables IGMP snooping for vlan HR
    }
    stp {
        disable;
    }
    rstp {
        disable;
    }
}
vlans {
    HR {
        vlan-id 100;
        l3-interface vlan.100;
    }
}
virtual-chassis {
    preprovisioned;
    member 0 {
        role routing-engine;
        serial-number BM0208105257;
    }
    member 1 {
        role routing-engine;
        serial-number BP0208137931;
    }
}
```

EX-VC-5 Configuration

```
#…truncated…
chassis {
    redundancy {
        graceful-switchover;
    }
    aggregated-devices {
        ethernet {
            device-count 2;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members SALES;
                }
            }
        }
    }
    ge-0/0/23 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members ENG;
                }
            }
        }
    }
    ge-2/0/0 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members ENG;
                }
            }
        }
    }
    ge-2/0/1 {
        unit 0 {
            family inet {
                address 172.18.16.34/30;
            }
        }
    }
    ge-2/0/23 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members SALES;
                }
            }
        }
```

```
}
ge-2/1/0 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae0;
    }
}
ge-2/1/1 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae0;
    }
}
ge-2/1/2 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae1;
    }
}
ge-2/1/3 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae1;
    }
}
ae0 {
    unit 0 {
        family inet {
            address 172.18.16.6/30;
        }
    }
}
ae1 {
    unit 0 {
        family inet {
            address 172.18.16.14/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.18.8.2/32;
        }
    }
}
vlan {
    unit 200 {
        family inet {
            address 172.18.10.1/24;
        }
    }
```

```
        unit 300 {
            family inet {
                address 172.18.11.1/24;
            }
        }
    }
    vme {
        unit 0 {
            family inet {
                address 172.19.59.195/24;
            }
        }
    }
}
routing-options {
    graceful-restart;
    router-id 172.18.8.2;
}
protocols {
    igmp;    # Explicitly enables IGMP
    ospf {
        area 0.0.0.1 {
            stub default-metric 10;
            interface ae0.0 {
                authentication {
                    simple-password "$9$PTF/uORlK8CtK8X7sYfTz3Ct"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface ae1.0 {
                authentication {
                    simple-password "$9$d4w2ajHmFnCZUnCtuEhVwYgZU"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface vme.0 {
                disable;
            }
            interface vlan.300 {
                passive;
            }
            interface lo0.0 {
                passive;
            }
            interface ge-2/0/1.0 {
                authentication {
                    simple-password "$9$ssgaUqmT/CujHCuO1yrYgoJjH"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface vlan.200 {
                passive;
            }
        }
```

```
    }
    pim {
        rp {  # Configuration of anycast-PIM on non-RP devices
            static {
                address 172.18.19.254;  # Uses the shared anycast address as the RP address
            }
        }
        interface all {
            mode sparse;  # Sets all interfaces to PIM sparse mode
        }
        interface vme.0 {
            disable;  # Explicitly disables PIM on the management interface
        }
    }
    igmp-snooping {  # Enables IGMP snooping for vlans SALES and ENG
        vlan SALES;
        vlan ENG;
    }
    stp {
        disable;
    }
    rstp {
        disable;
    }
}
vlans {
    ENG {
        vlan-id 200;
        l3-interface vlan.200;
    }
    SALES {
        vlan-id 300;
        l3-interface vlan.300;
    }
}
virtual-chassis {
    preprovisioned;
    /* ex-vc-5*/
    member 0 {
        role routing-engine;
        serial-number BP0208180059;
    }
    /* ex-vc-7 */
    member 1 {
        role routing-engine;
        serial-number BP0208180087;
    }
    /* ex-vc-6 */
    member 2 {
        role line-card;
        serial-number BQ0208189143;
    }
}
```

EX-VC-8 Configuration

```
#… truncated…
chassis {
    redundancy {
        graceful-switchover;
    }
    aggregated-devices {
        ethernet {
            device-count 2;
        }
    }
}
interfaces {
    ge-0/1/0 {
        ether-options {
            speed {
                1g;
            }
            802.3ad ae0;
        }
    }
    ge-0/1/1 {
        ether-options {
            speed {
                1g;
            }
            802.3ad ae1;
        }
    }
    ge-1/1/0 {
        ether-options {
            speed {
                1g;
            }
            802.3ad ae0;
        }
    }
    ge-1/1/1 {
        ether-options {
            speed {
                1g;
            }
            802.3ad ae1;
        }
    }
    ge-2/0/0 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members SUPPORT;
                }
            }
        }
    }
    ge-2/0/1 {
        unit 0 {
            family inet {
                address 172.18.16.98/30;
            }
        }
    }
    ge-2/0/23 {
        unit 0 {
```

```
            family ethernet-switching {
                port-mode access;
                vlan {
                    members SUPPORT;
                }
            }
        }
    }
    ae0 {
        unit 0 {
            family inet {
                address 172.18.16.66/30;
            }
        }
    }
    ae1 {
        unit 0 {
            family inet {
                address 172.18.16.74/30;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.18.12.1/32;
            }
        }
    }
    vlan {
        unit 400 {
            family inet {
                address 172.18.13.1/24;
            }
        }
    }
    vme {
        unit 0 {
            family inet {
                address 172.19.59.198/24;
            }
        }
    }
}
routing-options {
    graceful-restart;
    router-id 172.18.12.1;
}
protocols {
    igmp;  # Explicitly enables IGMP
    ospf {
        area 0.0.0.2 {
            stub default-metric 10;
            interface ae0.0 {
                authentication {
```

```
                        simple-password "$9$XwxNVYJGifT3goT369OBxNdw4a"; ## SECRET-DATA
                    }
                    bfd-liveness-detection {
                        minimum-interval 300;
                    }
                }
                interface ae1.0 {
                    authentication {
                        simple-password "$9$Ku5WX-YgJHqfVwqfTzCAvWL7bs"; ## SECRET-DATA
                    }
                    bfd-liveness-detection {
                        minimum-interval 300;
                    }
                }
                interface vlan.400 {
                    passive;
                }
                interface vme.0 {
                    disable;
                }
                interface lo0.0 {
                    passive;
                }
                interface ge-2/0/1.0 {
                    authentication {
                        simple-password "$9$ELeSlM7-waZj8XZjHqQzhSrKLx"; ## SECRET-DATA
                    }
                    bfd-liveness-detection {
                        minimum-interval 300;
                    }
                }
            }
        }
        pim {
            rp {  # Configuration of anycast-PIM on non-RP devices
                static {
                    address 172.18.19.254;  # Uses the shared anycast address as the RP address
                }
            }
            interface all {
                mode sparse;  # Sets all interfaces to PIM sparse mode
            }
            interface vme.0 {
                disable;  # Explicitly disables PIM on the management interface
            }
        }
        igmp-snooping {
            vlan SUPPORT;    # Enables IGMP snooping for vlan SUPPORT
        }
        stp {
            disable;
        }
        rstp {
            disable;
        }
    }
    vlans {
        SUPPORT {
            vlan-id 400;
            l3-interface vlan.400;
```

```
    }
}
virtual-chassis {
    preprovisioned;
    /* ex-vc-8 */
    member 0 {
        role routing-engine;
        serial-number BN0208189106;
    }
    /* ex-vc-9 */
    member 1 {
        role routing-engine;
        serial-number BP0208180160;
    }
    /* ex-vc-10 */
    member 2 {
        role line-card;
        serial-number BP0208180149;
    }
}
```

EX-FC-2 Configuration

```
#… truncated …
chassis {
    aggregated-devices {
        ethernet {
            device-count 2;
        }
    }
}
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching {
                vlan {
                    members FINANCE;
                }
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family inet {
                address 172.18.16.97/30;
            }
        }
    }
    ge-0/0/19 {
        unit 0 {
            family inet {
                address 172.18.16.101/30;
            }
        }
    }
    ge-0/1/0 {
        ether-options {
            speed {
```

```
            1g;
        }
        802.3ad ae0;
    }
}
ge-0/1/1 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae1;
    }
}
ge-0/1/2 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae0;
    }
}
ge-0/1/3 {
    ether-options {
        speed {
            1g;
        }
        802.3ad ae1;
    }
}
ae0 {
    unit 0 {
        family inet {
            address 172.18.16.70/30;
        }
    }
}
ae1 {
    unit 0 {
        family inet {
            address 172.18.16.78/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.18.12.2/32;
        }
    }
}
me0 {
    unit 0 {
        family inet {
            address 172.19.59.208/24;
        }
    }
}
vlan {
    unit 500 {
        family inet {
```

```
                address 172.18.14.1/24;
            }
        }
    }
}
routing-options {
    graceful-restart;
    router-id 172.18.12.2;
}
protocols {
    igmp;  # Explicitly enables IGMP
    ospf {
        area 0.0.0.2 {
            stub default-metric 10;
            interface ae0.0 {
                authentication {
                    simple-password "$9$U-iqf36A1RSTzRSreXxDikmQF"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface ae1.0 {
                authentication {
                    simple-password "$9$2roZjmfzCtOHqtO1RlegoJUk."; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface me0.0 {
                disable;
            }
            interface lo0.0 {
                passive;
            }
            interface vlan.500 {
                passive;
            }
            interface ge-0/0/1.0 {
                authentication {
                    simple-password "$9$Qbiw3/t1RSM87uO87-V4oz36C0B"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface ge-0/0/19.0 {
                passive;
            }
        }
    }
    pim {
        rp {  # Configuration of anycast-PIM on non-RP devices
            static {
                address 172.18.19.254;  # Uses the shared anycast address as the RP address
            }
        }
        interface all {
            mode sparse;  # Sets all interfaces to PIM sparse mode
```

```
        }
        interface me0.0 {
            disable;  # Explicitly disables PIM on the management interface
        }
    }
    igmp-snooping {
        vlan FINANCE;      # Enables IGMP snooping for vlan FINANCE
    }
    stp {
        disable;
    }
    rstp {
        disable;
    }
}
vlans {
    FINANCE {
        vlan-id 500;
        l3-interface vlan.500;
    }
}
```

MX-A Configuration

```
#… truncated…
chassis {
    redundancy {
        graceful-switchover;
    }
    aggregated-devices {
        ethernet {
            device-count 4;
        }
    }
    fpc 5 {     # Required to create tunnel interface to de-encapsulate PIM register messages
        pic 0 {
            tunnel-services {
                bandwidth 1g;
            }
        }
    }
}
interfaces {
    ge-5/0/0 {
        gigether-options {
            802.3ad ae0;
        }
    }
    ge-5/0/1 {
        gigether-options {
            802.3ad ae1;
        }
    }
    ge-5/0/2 {
        gigether-options {
            802.3ad ae2;
        }
    }
```

```
ge-5/0/3 {
    gigether-options {
        802.3ad ae3;
    }
}
ge-5/0/4 {
    gigether-options {
        802.3ad ae0;
    }
}
ge-5/0/5 {
    gigether-options {
        802.3ad ae1;
    }
}
ge-5/0/6 {
    gigether-options {
        802.3ad ae2;
    }
}
ge-5/0/7 {
    gigether-options {
        802.3ad ae3;
    }
}
ge-5/0/8 {
    unit 0 {
        family inet {
            address 172.18.16.129/30;
        }
    }
}
ge-5/0/9 {
    unit 0 {
        family inet {
            address 172.18.16.133/30;
        }
    }
}
ge-5/1/0 {
    unit 0 {
        family inet {
            address 172.18.16.141/30;
        }
    }
}
ae0 {
    unit 0 {
        family inet {
            address 172.18.16.1/30;
        }
    }
}
ae1 {
    unit 0 {
        family inet {
            address 172.18.16.5/30;
        }
    }
}
```

```
    ae2 {
        unit 0 {
            family inet {
                address 172.18.16.65/30;
            }
        }
    }
    ae3 {
        unit 0 {
            family inet {
                address 172.18.16.69/30;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.18.19.1/32 {
                    primary;
                }
                address 172.18.19.254/32; # Loopback address used as shared anycast address
            }
        }
    }
}
routing-options {
    graceful-restart;
    static {
        route 200.0.1.0/24 reject;
        route 200.0.2.0/24 reject;
    }
    router-id 172.18.19.1;
}
protocols {
    ospf {
        export stat;
        area 0.0.0.1 {
            stub default-metric 10;
            area-range 172.18.8.0/22;
            area-range 172.18.16.0/26;
            interface ae0.0 {
                authentication {
                    simple-password "$9$0fRe1EyM87s2alK2aZU.mO1RhlK"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface ae1.0 {
                authentication {
                    simple-password "$9$-PbYoDi.z39JG39ApREdbs2JG"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
        }
        area 0.0.0.2 {
            stub default-metric 10;
            area-range 172.18.12.0/22;
```

```
            area-range 172.18.16.64/26;
            interface ae2.0 {
                authentication {
                    simple-password "$9$1dIESeLxdgoGvWoGDif5IEcrM8"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface ae3.0 {
                authentication {
                    simple-password "$9$5znCO1hKMXtuMX7-2gTz3/p0"; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
        }
        area 0.0.0.0 {
            interface ge-5/0/8.0 {
                authentication {
                    md5 10 key "$9$zNHSn9pIEyWLN0BLNdboaFn/9uO"; ## SECRET-DATA
                }
            }
            interface ge-5/0/9.0 {
                authentication {
                    md5 10 key "$9$aPGjqTz6uORmfORhSMWJGDj.P"; ## SECRET-DATA
                }
            }
            interface lo0.0 {
                passive;
            }
            interface fxp0.0 {
                disable;
            }
            interface ge-5/1/0.0 {
                passive;
            }
        }
    }
    pim {
        rp {  # Configuration of anycast-PIM on RP devices
            local {
                family inet {
                    address 172.18.19.254;  # Anycast-PIM address shared by the RPs
                    anycast-pim {
                        rp-set {
                            address 172.18.19.2;  # Local address of the other RP
                        }
                        local-address 172.18.19.1;   # Local address of this RP
                    }
                }
            }
        }
        interface all {
            mode sparse;  # Sets all interfaces to PIM sparse mode
        }
        interface fxp0.0 {
            disable;  # Explicitly disables PIM on the management interface
        }
```

```
    }
}
policy-options {
    policy-statement stat {
        from protocol static;
        then accept;
    }
}
```

MX-B Configuration

```
#… truncated…
chassis {
    redundancy {
        graceful-switchover;
    }
    aggregated-devices {
        ethernet {
            device-count 4;
        }
    }
    fpc 2 {     # Required to create tunnel interface to de-encapsulate PIM register messages
        pic 0 {
            tunnel-services {
                bandwidth 1g;
            }
        }
    }
}
interfaces {
    ge-2/0/0 {
        gigether-options {
            802.3ad ae0;
        }
    }
    ge-2/0/1 {
        gigether-options {
            802.3ad ae1;
        }
    }
    ge-2/0/2 {
        gigether-options {
            802.3ad ae2;
        }
    }
    ge-2/0/3 {
        gigether-options {
            802.3ad ae3;
        }
    }
    ge-2/0/4 {
        gigether-options {
            802.3ad ae0;
        }
    }
    ge-2/0/5 {
        gigether-options {
            802.3ad ae1;
```

```
    }
}
ge-2/0/6 {
    gigether-options {
        802.3ad ae2;
    }
}
ge-2/0/7 {
    gigether-options {
        802.3ad ae3;
    }
}
ge-2/0/8 {
    unit 0 {
        family inet {
            address 172.18.16.130/30;
        }
    }
}
ge-2/0/9 {
    unit 0 {
        family inet {
            address 172.18.16.137/30;
        }
    }
}
ae0 {
    unit 0 {
        family inet {
            address 172.18.16.9/30;
        }
    }
}
ae1 {
    unit 0 {
        family inet {
            address 172.18.16.13/30;
        }
    }
}
ae2 {
    unit 0 {
        family inet {
            address 172.18.16.73/30;
        }
    }
}
ae3 {
    unit 0 {
        family inet {
            address 172.18.16.77/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.18.19.2/32 {
                primary;
            }
            address 172.18.19.254/32;  # Loopback address used as shared anycast address
```

```
            }
        }
    }
}
routing-options {
    graceful-restart;
    static {
        route 200.0.1.0/24 reject;
        route 200.0.2.0/24 reject;
    }
    router-id 172.18.19.2;
}
protocols {
    ospf {
        export stat;
        area 0.0.0.1 {
            stub default-metric 10;
            area-range 172.18.8.0/22;
            area-range 172.18.16.0/26;
            interface ae0.0 {
                authentication {
                    simple-password “$9$d6w2ajHmFnCZUnCtuEhVwYgZU”; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface ae1.0 {
                authentication {
                    simple-password “$9$AXE/uBElK8db2cyb24aiHtuO1cy”; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
        }
        area 0.0.0.2 {
            stub default-metric 10;
            area-range 172.18.12.0/22;
            area-range 172.18.16.64/26;
            interface ae2.0 {
                authentication {
                    simple-password “$9$m5z6p0IreW9AeWLxwsP5Q3Ct”; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
            interface ae3.0 {
                authentication {
                    simple-password “$9$hRNyeWNdsJGiLxGik.zFcylvX7”; ## SECRET-DATA
                }
                bfd-liveness-detection {
                    minimum-interval 300;
                }
            }
        }
        area 0.0.0.0 {
            interface ge-2/0/8.0 {
                authentication {
                    md5 10 key “$9$MKTL7VgoGqmTwYmTz3tpWLx7bs”; ## SECRET-DATA
                }
            }
```

```
            interface ge-2/0/9.0 {
                authentication {
                    md5 10 key "$9$m5z6p0IreW9AeWLxwsP5Qz/C"; ## SECRET-DATA
                }
            }
            interface lo0.0 {
                passive;
            }
            interface fxp0.0 {
                disable;
            }
        }
    }
    pim {
        traceoptions {   # Traceoptions for PIM troubleshooting
            file pim;
            flag join detail;
            flag task detail;
            flag state detail;
            flag packets detail;
        }
        rp {  # Uses the shared anycast address as the RPs address
            local {
                family inet {
                    address 172.18.19.254;    # Anycast-PIM address shared by the RPs
                    anycast-pim {
                        rp-set {
                            address 172.18.19.1;  # Local address of the other RP
                        }
                        local-address 172.18.19.2;    # Local address of this RP
                    }
                }
            }
        }
        interface all {
            mode sparse;  # Sets all interfaces to PIM sparse mode
        }
        interface fxp0.0 {
            disable;  # Explicitly disables PIM on the management interface
        }
    }
}
policy-options {
    policy-statement stat {
        from protocol static;
        then accept;
    }
}
```

Juniper Networks M120 Multiservice Edge Router Configuration

```
#… truncated…
chassis {
    redundancy {
        graceful-switchover;
    }
}
interfaces {
    ge-5/0/0 {
        unit 0 {
            family inet {
```

```
                address 172.18.16.134/30;
            }
        }
    }
    ge-5/0/1 {
        unit 0 {
            family inet {
                address 172.18.16.138/30;
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 172.18.19.3/32 {
                    primary;
                }
            }
        }
    }
}
routing-options {
    graceful-restart;
    router-id 172.18.19.3;
}
protocols {
    ospf {
        area 0.0.0.0 {
            interface ge-5/0/0.0 {
                authentication {
                    md5 10 key "$9$aLGjqTz6uORmfORhSMWJGDj.P"; ## SECRET-DATA
                }
            }
            interface ge-5/0/1.0 {
                authentication {
                    md5 10 key "$9$.fQntu1ylM/ClM8XbwmfTQ69"; ## SECRET-DATA
                }
            }
            interface lo0.0 {
                passive;
            }
            interface fxp0.0 {
                disable;
            }
        }
    }
    pim {
        rp {  # Configuration of anycast-PIM on non-RP devices
            static {
                address 172.18.19.254;  # Uses the shared anycast address as the RP address
            }
        }
        interface all {
            mode sparse;  # Sets all interfaces to PIM sparse mode
        }
        interface fxp0.0 {
            disable;  # Explicitly disables PIM on the management interface
        }
    }
}
```

Verification

This section lists some commands that can be used to verify the multicast setup along with sample outputs:

General Multicast

Show Multicast RPF

This command displays information about the PIM RPF table similar to what can be seen in the routing table. The source prefix, the protocol, the upstream interface and the neighbor that it was learned from are all displayed:

```
lab@ex-vc-1> show multicast rpf
Multicast RPF table: inet.0 , 33 entries

0.0.0.0/0
    Protocol: OSPF
    Interface: ae1.0
    Neighbor: 172.18.16.9

172.18.8.1/32
    Protocol: Direct
    Interface: lo0.0

#…truncated…

224.0.0.2/32
    Protocol: PIM

224.0.0.5/32
    Protocol: OSPF

224.0.0.13/32
    Protocol: PIM

224.0.0.22/32
    Protocol: IGMP

Multicast RPF table: inet6.0 , 2 entries

ff02::2/128
    Protocol: PIM

ff02::d/128
    Protocol: PIM

{master:0}
```

Show Multicast Route Extensive

This command can be used to examine the multicast routing table that acts as a cache and is based on both IGMP group activity and RPF information. Note “show route table inet.1 extensive” can also be used to examine this table.

```
lab@ex-vc-1> show multicast route extensive

Group: 225.0.0.1
    Source: 172.18.16.102/32
    Upstream interface: ae0.0
    Downstream interface list:
        vlan.100
    Session description: MALLOC
    Statistics: 0 kBps, 0 pps, 3209 packets
    Next-hop ID: 131070
```

```
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 10

Group: 225.0.0.1
    Source: 172.18.16.142/32
    Upstream interface: ae0.0
    Downstream interface list:
        vlan.100
    Session description: MALLOC
    Statistics: 0 kBps, 0 pps, 3184 packets
    Next-hop ID: 131070
    Upstream protocol: PIM
    Route state: Active
    Forwarding state: Forwarding
    Cache lifetime/timeout: 360 seconds
    Wrong incoming interface notifications: 1

{master:0}
```

Show Multicast Usage

```
lab@ex-vc-1> show multicast usage
Group           Sources Packets              Bytes
225.0.0.1       2       6403                 53998


Prefix          /len Groups Packets              Bytes
172.18.16.102   /32  1      3214                 27574
172.18.16.142   /32  1      3189                 26424

{master:0}
```

Show Multicast Next-Hops

```
lab@ex-vc-1> show multicast next-hops
Family: INET
ID          Refcount KRefcount Downstream interface
131070             4         2 vlan.100

Family: INET6

{master:0}
```

PIM

Show PIM RPS Extensive

This command displays information about RPs, how they were learned, the group range they handle, and the ones that are currently active. In the example below, since this is the RP, the output also includes anycast PIM RP information. The configuration of a tunnel PIC on this MX Series device resulted in the creation of a de-encapsulation interface that allows the RP to receive multicast traffic from the source. This interface is indicated by pd-5/0/10.32769 in the example below.

```
lab@Madras-MX-A> show pim rps extensive
Instance: PIM.master
Address family INET

RP: 172.18.19.254
Learned via: static configuration
Time Active: 1w2d 16:09:14
Holdtime: 0
Device Index: 153
Subunit: 32769
Interface: pd-5/0/10.32769
Group Ranges:
        224.0.0.0/4
Register State for RP:
Group           Source          FirstHop        RP Address      State    Timeout
225.0.0.1       172.18.16.102   172.18.19.2     172.18.19.254   Receive      261
225.0.0.2       172.18.16.102   172.18.19.2     172.18.19.254   Receive      281
225.0.0.3       172.18.16.102   172.18.19.2     172.18.19.254   Receive        0
Anycast PIM rpset:
        172.18.19.2
Anycast PIM local address used: 172.18.19.1
Anycast PIM Register State:
Group                          Source                         Origin
225.0.0.1                      172.18.16.142                  DIRECT
225.0.0.2                      172.18.16.142                  DIRECT
225.0.0.3                      172.18.16.142                  DIRECT

Address family INET6

{master}
```

Show PIM Interface

Use this command to list the currently configured and operational interfaces for PIM. As shown in the example below, this includes the name of each interface, its operational state (up/down), its mode (sparse, dense, or sparse-dense), the IP version supported (IPv4 in this case), the PIM version (2 by default), the current state of the interface, the neighbor count, the join count, and the designated router address.

If the interface is responsible for forwarding traffic, the state shows as “DR”. If it is not responsible for forwarding traffic, it shows as “NotDR”. Note that this command also shows the de-encapsulation interface (pd-5/0/10.32769) which is a non-broadcast interface as indicated by its state (point to point).

```
lab@Madras-MX-A> show pim interfaces
Instance: PIM.master

Name                Stat      Mode       IP  V    State              NbrCnt    JoinCnt   DR address
ae0.0               Up        Sparse      4  2    NotDR              1         0         172.18.16.2
ae1.0               Up        Sparse      4  2    NotDR              1         0         172.18.16.6
ae2.0               Up        Sparse      4  2    NotDR              1         1         172.18.16.66
ae3.0               Up        Sparse      4  2    NotDR              1         3         172.18.16.70
ge-5/0/8.0          Up        Sparse      4  2    NotDR              1         0         172.18.16.130
ge-5/0/9.0          Up        Sparse      4  2    NotDR              1         0         172.18.16.134
ge-5/1/0.0          Up        Sparse      4  2    DR                           0         0         172.18.16.141
lo0.0               Up        Sparse      4  2    DR                           0         0         172.18.19.1
pd-5/0/10.32769     Up    Sparse      4  2    P2P                  0         0
{master}
```

Show PIM Neighbor

Use this command to display information about neighboring routers running PIM as shown in the example below:

```
lab@Madras-MX-A> show pim neighbors
Instance: PIM.master
B = Bidirectional Capable, G = Generation Identifier,
H = Hello Option Holdtime, L = Hello Option LAN Prune Delay,
P = Hello Option DR Priority

Interface           IP   V Mode        Option              Uptime              Neighbor addr
ae0.0                    4    2                  HPLG                1w2d16h   172.18.16.2
ae1.0                    4    2                  HPLG                1w2d16h   172.18.16.6
ae2.0                    4    2                  HPLG                1w2d16h   172.18.16.66
ae3.0                    4    2                  HPLG                1w2d16h   172.18.16.70
ge-5/0/8.0               4    2                  HPLG                1w2d16h   172.18.16.130
ge-5/0/9.0               4    2                  HPLG                1w2d16h   172.18.16.134

{master}
```

Show PIM Join

Use this command to display PIM join states. In the example below, the use of the asterisk (*) for the source in the first entry indicates a shared rendezvous-point tree, which is a (*,G) state. The second and third entries indicate the use of the shortest-path tree, which is a (S,G) state. The command also displays the upstream interface for each entry.

```
lab@ex-vc-1> show pim join
Instance: PIM.master Family: INET
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

Group: 225.0.0.1
    Source: *
    RP: 172.18.19.254
    Flags: sparse,rptree,wildcard
    Upstream interface: ae1.0

Group: 225.0.0.1
    Source: 172.18.16.102
    Flags: sparse,spt
    Upstream interface: ae0.0

Group: 225.0.0.1
    Source: 172.18.16.142
    Flags: sparse,spt
    Upstream interface: ae0.0

Instance: PIM.master Family: INET6
R = Rendezvous Point Tree, S = Sparse, W = Wildcard

{master:0}
```

Show PIM Source Detail

Use this command to display information about the active sources and what the RPF interface is for each source. In the example below, the source of the traffic is 172.18.19.254 since it is the shared anycast-PIM RP address.

```
lab@ex-vc-1> show pim source detail
Instance: PIM.master Family: INET
```

```
Source 172.18.19.254
    Prefix 172.18.19.254/32
    Upstream interface ae1.0
    Upstream neighbor 172.18.16.9
    Active groups:225.0.0.1

Instance: PIM.master Family: INET6

{master:0}
```

Show PIM Statistics

This command displays PIM-related messages and error counts as shown in the example below:

```
lab@Madras-MX-A> show pim statistics

PIM Message type        Received       Sent  Rx errors
V2 Hello                  175925     205260          0
V2 Register                41580      41561          0
V2 Register Stop           41542      41536          0
V2 Join Prune              41655      13863          0
V2 Bootstrap                   0          0          0
V2 Assert                      0          0          0
V2 Graft                       0          0          0
V2 Graft Ack                   0          0          0
V2 Candidate RP                0          0          0
V1 Query                       0          0          0
V1 Register                    0          0          0
V1 Register Stop               0          0          0
V1 Join Prune                  0          0          0
V1 RP Reachability             0          0          0
V1 Assert                      0          0          0
V1 Graft                       0          0          0
V1 Graft Ack                   0          0          0
AutoRP Announce                0          0          0
AutoRP Mapping                 0          0          0
AutoRP Unknown type            0
Anycast Register           41526      41561          0
Anycast Register Stop      41542      41529          0

Global Statistics

Hello dropped on neighbor policy         0
Unknown type                             0
V1 Unknown type                          0
Unknown Version                          0
Neighbor unknown                         0
Bad Length                               0
Bad Checksum                             0
Bad Receive If                           0
Rx Bad Data                              0
Rx Intf disabled                         0
Rx V1 Require V2                         0
Rx V2 Require V1                         0
Rx Register not RP                       0
Rx Register no route                     0
Rx Register no decap if                  0
Null Register Timeout                    0
RP Filtered Source                       0
Rx Unknown Reg Stop                      0
Rx Join/Prune no state                   0
Rx Join/Prune on upstream if             0
Rx sparse join for dense group           0
Rx Graft/Graft Ack no state              0
```

```
Rx Graft on upstream if                  0
Rx CRP not BSR                           0
Rx BSR when BSR                          0
Rx BSR not RPF if                        0
Rx unknown hello opt                     0
Rx data no state                         0
Rx RP no state                           0
Rx aggregate                             0
Rx malformed packet                      0
No RP                                    0
No register encap if                     0
No route upstream                        0
Nexthop Unusable                         0
RP mismatch                              0
RPF neighbor unknown                     0
Rx Joins/Prunes filtered                 0
Embedded-RP invalid addr                 0
Embedded-RP limit exceed                 0
Embedded-RP added                        0
Embedded-RP removed                      0
Rx Register msgs filtering drop          0
Tx Register msgs filtering drop          0

{master}
```

IGMP

Show IGMP Interface

This command lists information about the interfaces running IGMP including their state, the IGMP querier address, the IGMP version, and the number of groups currently active. It also lists the derived and configured timer and counter parameters for IGMP (all use the default values in the example below).

```
lab@ex-vc-1> show igmp interface
Interface: ae0.0
    Querier: 172.18.16.1
    State:         Up Timeout:     166 Version:  2 Groups:      0
    Immediate leave: Off
    Promiscuous mode: Off
Interface: ae1.0
    Querier: 172.18.16.9
    State:         Up Timeout:     167 Version:  2 Groups:      0
    Immediate leave: Off
    Promiscuous mode: Off
Interface: vlan.100
    Querier: 172.18.9.1
    State:         Up Timeout:    None Version:  2 Groups:      1
    Immediate leave: Off
    Promiscuous mode: Off
Interface: ge-0/0/2.0
    Querier: 172.18.16.33
    State:         Up Timeout:    None Version:  2 Groups:      0
    Immediate leave: Off
    Promiscuous mode: Off
```

```
Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

{master:0}
```

Show IGMP Group

Use this command to show the groups joined by directly connected hosts and other routers. Note that if IGMP is explicitly enabled but PIM is not enabled on an interface, the status of that interface is shown as "Up" but is omitted from the output of the "show igmp group".

```
lab@ex-vc-1> show igmp group
Interface: vlan.100
    Group: 225.0.0.1
        Source: 0.0.0.0
        Last reported by: 172.18.9.2
        Timeout:     187 Type: Dynamic
Interface: local
    Group: 224.0.0.2
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.5
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.6
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic
    Group: 224.0.0.22
        Source: 0.0.0.0
        Last reported by: Local
        Timeout:       0 Type: Dynamic

{master:0}
```

Show IGMP Statistics

This command displays IGMP-related messages and error counts as shown in the example below:

```
lab@ex-vc-1> show igmp statistics
IGMP packet statistics for all interfaces
IGMP Message type      Received       Sent  Rx errors
Membership Query          13425      13436          0
V1 Membership Report          0          0          0
DVMRP                         0          0          0
PIM V1                        0          0          0
Cisco Trace                   0          0          0
V2 Membership Report       7335          0          0
Group Leave                   0          0          0
Mtrace Response               0          0          0
Mtrace Request                0          0          0
```

```
Domain Wide Report            0          0          0
V3 Membership Report          0          0          0
Other Unknown types                                 0
IGMP v3 unsupported type                            0
IGMP v3 source required for SSM                     0
IGMP v3 mode not applicable for SSM                 0

IGMP Global Statistics
Bad Length                    0
Bad Checksum                  0
Bad Receive If                0
Rx non-local                  0
Timed out                   619
Rejected Report               0
Total Interfaces              4

{master:0}
{master:0}
```

Show IGMP Snooping Membership

The IGMP snooping cache can be displayed through “show igmp-snooping membership” command as shown below:

```
lab@ex-vc-1> show igmp-snooping membership
VLAN: HR
    225.0.0.1      *               171 secs
        Interfaces: ge-0/0/23.0

{master:0}
```

Show IGMP Snooping Route Ethernet Switching

The forwarding state that is generated can be seen through “show igmp-snooping route ethernet-switching”:

```
lab@ex-vc-1> show igmp-snooping route ethernet-switching
VLAN             Group          Next-hop
HR               224.0.0.0, *
HR               225.0.0.1, *       1309

{master:0}
```

Show IGMP Snooping Statistics

This command shows message and error counts related to IGMP snooping:

```
lab@ex-vc-1> show igmp-snooping statistics
Bad length: 0 Bad checksum: 0 Invalid interface: 0
Not local: 0 Receive unknown: 0 Timed out: 0


    IGMP Type           Received      Transmited    Recv Errors
    Queries:                 0               2              0
    Reports:              6716               0              0
    Leaves:                  0               0              0
    Other:                   0               0              0

{master:0}
```

Show IGMP Snooping VLANS

```
lab@ex-vc-1> show igmp-snooping vlans detail
VLAN: HR, Tag: 100, vlan-interface: vlan.100
Membership timeout: 260, Querier timeout: 255
        Interface: ge-0/0/0.0, untagged, Groups: 0, Reporters: 0
        Interface: ge-0/0/23.0, untagged, Groups: 1, Reporters: 1

{master:0}
```

Troubleshooting

Confirming the Presence of De-Encapsulation (pd) or Encapsulation (pe) Interfaces

- Show interfaces terse | match “pd|pe”

Verifying Multicast Routes in Forwarding Table

- Show route forwarding-table multicast destination <prefix>
- Request pfe execute target <fpc> command “sh nhdb id <id> ext”

Clearing Statistics and Usage

- Clear pim join
- Clear multicast statistics
- Clear multicast usage
- Clear IGMP statistics
- Clear IGMP snooping membership
- Clear IGMP snooping statistics

Trace Options

- Set protocols pim traceoptions flag hello detail
- Set protocols pim traceoptions flag join detail
- Set protocols pim traceoptions flag prune detail
- Set protocols pim traceoptions flag packets detail
- Set protocols pim traceoptions flag state detail
- Set protocols pim traceoptions flag task detail
- Set protocols igmp traceoptions flag query detail
- Set protocols igmp traceoptions flag report detail
- Set protocols igmp traceoptions flag leave detail (IGMP version 2 only)
- Set protocols igmp traceoptions flag packets detail
- Set protocols igmp-snooping traceoptions flag query detail
- Set protocols igmp-snooping traceoptions flag report detail
- Set protocols igmp-snooping traceoptions flag leave detail (IGMP version 2 only)
- Set protocols igmp-snooping traceoptions flag packets detail

Summary

Juniper Networks provides enterprise customers with a variety of solutions to implement multicast forwarding based on their specific requirements to deliver applications such as interactive distance learning, corporate video conferencing, inventory updates, software, and content distribution. This document provided a brief overview of some of the options that Junos OS offers and presented the pros and cons of each. It then presented implementation guidelines for IP multicast deployments in enterprise networks using the EX Series Ethernet Switches and MX Series 3D Universal Edge Routers and two core multicast protocols: PIM sparse mode and IGMPv2. Finally, an implementation example was presented with detailed configurations, verifications, and troubleshooting procedures.

Appendixes

Appendix A: Conventions/Glossary

BSR	Bootstrap router
Candidate RP	Candidate RP
DPC	Dense Port Concentrator
DVMRP	Distance Vector Multicast Routing Protocol
FPC	Flexible PIC Concentrator
IGMP	Internet Group Management Protocol
MSDP	Multicast Source Discovery Protocol
OSPF	Open Shortest Path First
PIC	Physical Interface Card
PIM	Protocol Independent Multicast
PIM DM	PIM dense mode
PIM SM	PIM sparse mode
PIM SSM	PIM source-specific multicast
P2P	Point to point
RE	Routing Engine
RP	Rendezvous point
RPF	Reverse path forwarding
SFP	Small form-factor pluggable transceiver
SPT	Shortest-path tree
VC	Virtual Chassis
VCP	Virtual Chassis Port
VLAN	Virtual LAN

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

8010062-001-EN Dec 2009