BROCADE

CONFIGURATION GUIDE

Brocade Network OS IP Multicast Configuration Guide, 7.2.0

Supporting Network OS 7.2.0

53-1005137-01
10 July 2017

Contents

Preface

Document conventions

The document conventions describe text formatting conventions, command syntax conventions, and important notice formats used in Brocade technical documentation.

Notes, cautions, and warnings

Notes, cautions, and warning statements may be used in this document. They are listed in the order of increasing severity of potential hazards.

NOTE
A Note provides a tip, guidance, or advice, emphasizes important information, or provides a reference to related information.

ATTENTION
An Attention statement indicates a stronger note, for example, to alert you when traffic might be interrupted or the device might reboot.

CAUTION
A Caution statement alerts you to situations that can be potentially hazardous to you or cause damage to hardware, firmware, software, or data.

DANGER
A Danger statement indicates conditions or situations that can be potentially lethal or extremely hazardous to you. Safety labels are also attached directly to products to warn of these conditions or situations.

Text formatting conventions

Text formatting conventions such as boldface, italic, or Courier font may be used to highlight specific words or phrases.

Format	Description
bold text	Identifies command names.
	Identifies keywords and operands.
	Identifies the names of GUI elements.
	Identifies text to enter in the GUI.
italic text	Identifies emphasis.
	Identifies variables.
	Identifies document titles.
`Courier font`	Identifies CLI output.

Format	Description
	Identifies command syntax examples.

Command syntax conventions

Bold and italic text identify command syntax components. Delimiters and operators define groupings of parameters and their logical relationships.

Convention	Description
bold text	Identifies command names, keywords, and command options.
italic text	Identifies a variable.
value	In Fibre Channel products, a fixed value provided as input to a command option is printed in plain text, for example, **--show** WWN.
[ ]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ **x** | **y** | **z** }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options. In Fibre Channel products, square brackets may be used instead for this purpose.
x | **y**	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, for example, passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, *member*[*member*...].
\	Indicates a "soft" line break in command examples. If a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Brocade resources

Visit the Brocade website to locate related documentation for your product and additional Brocade resources.

White papers, data sheets, and the most recent versions of Brocade software and hardware manuals are available at www.brocade.com. Product documentation for all supported releases is available to registered users at MyBrocade.

Click the **Support** tab and select **Document Library** to access product documentation on MyBrocade or www.brocade.com. You can locate documentation by product or by operating system.

Release notes are bundled with software downloads on MyBrocade. Links to software downloads are available on the MyBrocade landing page and in the Document Library.

Document feedback

Quality is our first concern at Brocade, and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. You can provide feedback in two ways:

- Through the online feedback form in the HTML documents posted on www.brocade.com
- By sending your feedback to documentation@brocade.com

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Contacting Brocade Technical Support

As a Brocade customer, you can contact Brocade Technical Support 24x7 online or by telephone. Brocade OEM customers should contact their OEM/solution provider.

Brocade customers

For product support information and the latest information on contacting the Technical Assistance Center, go to www.brocade.com and select **Support**.

If you have purchased Brocade product support directly from Brocade, use one of the following methods to contact the Brocade Technical Assistance Center 24x7.

Online	Telephone
Preferred method of contact for non-urgent issues: • Case management through the MyBrocade portal. • Quick Access links to Knowledge Base, Community, Document Library, Software Downloads and Licensing tools	Required for Sev 1-Critical and Sev 2-High issues: • Continental US: 1-800-752-8061 • Europe, Middle East, Africa, and Asia Pacific: +800-AT FIBREE (+800 28 34 27 33) • Toll-free numbers are available in many countries. • For areas unable to access a toll-free number: +1-408-333-6061

Brocade OEM customers

If you have purchased Brocade product support from a Brocade OEM/solution provider, contact your OEM/solution provider for all of your product support needs.

- OEM/solution providers are trained and certified by Brocade to support Brocade® products.
- Brocade provides backline support for issues that cannot be resolved by the OEM/solution provider.
- Brocade Supplemental Support augments your existing OEM support contract, providing direct access to Brocade expertise. For more information, contact Brocade or your OEM.
- For questions regarding service levels and response times, contact your OEM/solution provider.

About This Document

Supported hardware and software

In those instances in which procedures or parts of procedures documented here apply to some switches but not to others, this guide identifies exactly which switches are supported and which are not.

Although many different software and hardware configurations are tested and supported by Brocade Communications Systems, Inc. for Network OS, documenting all possible configurations and scenarios is beyond the scope of this document.

The following hardware platforms are supported by this release of Network OS:

- Brocade VDX 2741
- Brocade VDX 2746
- Brocade VDX 6740
  - Brocade VDX 6740-48
  - Brocade VDX 6740-64
- Brocade VDX 6740T
  - Brocade VDX 6740T-48
  - Brocade VDX 6740T-64
  - Brocade VDX 6740T-1G
- Brocade VDX 6940-36Q
- Brocade VDX 6940-144S
- Brocade VDX 8770
  - Brocade VDX 8770-4
  - Brocade VDX 8770-8

To obtain information about a Network OS version other than this release, refer to the documentation specific to that version.

Using the Network OS CLI

For complete instructions and support for using the Brocade Network OS command line interface (CLI), refer to the *Brocade Network OS Command Reference*.

What's new in this document

This document supports the following features introduced in Network OS 7.2.0:

- Enhancement to the IGMP snooping querier feature

For complete information, refer to the *Network OS Release Notes*.

IP Multicast

IP multicast overview

Multicast protocols allow a group or channel to be accessed over different networks by multiple stations (clients) for the receipt and transmission of multicast data. Distribution of stock quotes, video transmissions such as news services and remote classrooms, and video conferencing are all examples of applications that use multicast routing.

Brocade devices support the Protocol-Independent Multicast (PIM) protocol, Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD).

PIM is a broadcast and pruning multicast protocol that delivers IP multicast datagrams. This protocol employs reverse path lookup check and pruning to allow source-specific multicast delivery trees to reach all group members. PIM builds a different multicast tree for each source and destination host group.

Multicast Listener Discovery (MLD) snooping is a multicast-constraining mechanism that runs on Layer 2 or Layer 3 devices to manage and control IPv6 multicast groups.

Multicast control packet and data forwarding through a Layer 2 switch is achieved by Layer 2 forwarding of the received multicast packets on all the member ports of the VLAN interfaces. This approach though simple is not bandwidth efficient, because only a subset of member ports may be connected to devices interested in receiving these multicast packets. In a worst-case scenario, the data gets forwarded to all port members of a VLAN with a large number of member ports even if only a single VLAN member is interested in receiving the data. Such scenarios can lead to loss of throughput for a switch upon receiving a high rate of multicast data traffic.

IGMP and MLD provide the functionality to save bandwidth and throughput by forwarding traffic to only interested receivers instead of all the member ports of the VLAN. IGMP snooping provides the specification for IPv4 and MLD snooping provides the specification for IPv6 data traffic forwarding.

MLD snooping is a multicast constraining mechanism that runs on Layer 2 or Layer 3 devices to manage and control IPv6 multicast groups. MLD snooping provides similar functionality for IPv6 as IGMP snooping for IPv4 by sending IPv6 multicast traffic only to the interested listeners. By listening to and analyzing MLD messages, a Layer 2 device running MLD snooping establishes mappings between ports and multicast MAC addresses or multicast IP addresses and forwards multicast data.

IP multicast message types

Multicast routers use IGMP or MLD to learn which groups have interested listeners on each of their attached physical networks. In any given subnet, one multicast router is elected to act as an IGMP or MLD querier.

The IGMP or MLD querier sends out the following types of queries to hosts:

- General query: Asks whether any host is listening to any group.
- Group-specific query: Asks whether any host is listening to a specific multicast group. This query is sent in response to a host leaving the multicast group and allows the router to quickly determine if any remaining hosts are interested in the group.

Hosts that are multicast listeners send the following kinds of messages:

- Membership report: Indicates that the host wants to join a particular multicast group.
- Leave report: Indicates that the host wants to leave a particular multicast group.

IPv4 Multicast Traffic Reduction

IGMP snooping overview

The forwarding of multicast control packets and data through a Layer 2 device configured with VLANs is most easily achieved by the Layer 2 forwarding of received multicast packets on all the member ports of the VLAN interfaces. However, this simple approach is not bandwidth efficient, because only a subset of member ports may be connected to devices interested in receiving those multicast packets. In a worst-case scenario, the data would get forwarded to all port members of a VLAN with a large number of member ports, even if only a single VLAN member is interested in receiving the data. Such scenarios can lead to loss of throughput for a device that gets hit by a high rate of multicast data traffic.

Internet Group Management Protocol (IGMP) snooping is a mechanism by which a Layer 2 device can effectively address this issue of inefficient multicast forwarding to VLAN port members. Snooping involves "learning" forwarding states for multicast data traffic on VLAN port members from the IGMP control (join/leave) packets received on them. The Layer 2 device also provides for a way to configure forwarding states statically through the CLI.

Multicast routing and IGMP snooping

Multicast routers use IGMP snooping to learn which groups have members on each of their attached physical networks. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership.

NOTE

"Multicast group memberships" means that at least one member of a multicast group on a given attached network is available.

There are two ways that hosts join multicast routing groups:

- By sending an unsolicited IGMP join request.
- By sending an IGMP join request as a response to a general query from a multicast router.

In response to the request, the device creates an entry in its Layer 2 forwarding table for that VLAN. When other hosts send join requests for the same multicast, the device adds them to the existing table entry. Only one entry is created per VLAN in the Layer 2 forwarding table for each multicast group.

vLAG and LAG primary port with IGMP snooping

The current data center Ethernet (DCE) implementation of vLAGs and LAGs uses the concept of a so-called primary port. One of the member ports of the vLAG and LAG is selected to be the primary port, and all multicast traffic egressing from the LAG or vLAG is sent on the primary port. Thus, normal hash-based forwarding is not performed for multicast traffic, whether it is control traffic or data. Now, consider the case where RBridge R1 receives an IGMP join request for group G1 on Po10, shown in the figure below. This causes Po10 to be added to the list of IGMP receivers for group G1. Now, assume that the primary port of the vLAG is the link connecting R4 and S1. Therefore, any multicast traffic received by the cluster for group G1 egresses on vLAG Po10 from R4 and not from R1, even though the original join was received on R1.

If the primary port for the vLAG changes, such as if the link between R4 and S1 in the figure below went down, then multicast traffic would egress out of the new primary port on the vLAG. In the above case, the new primary port would be the link connecting R1 and S1.

FIGURE 1 IGMP snooping in Brocade VCS Fabric mode

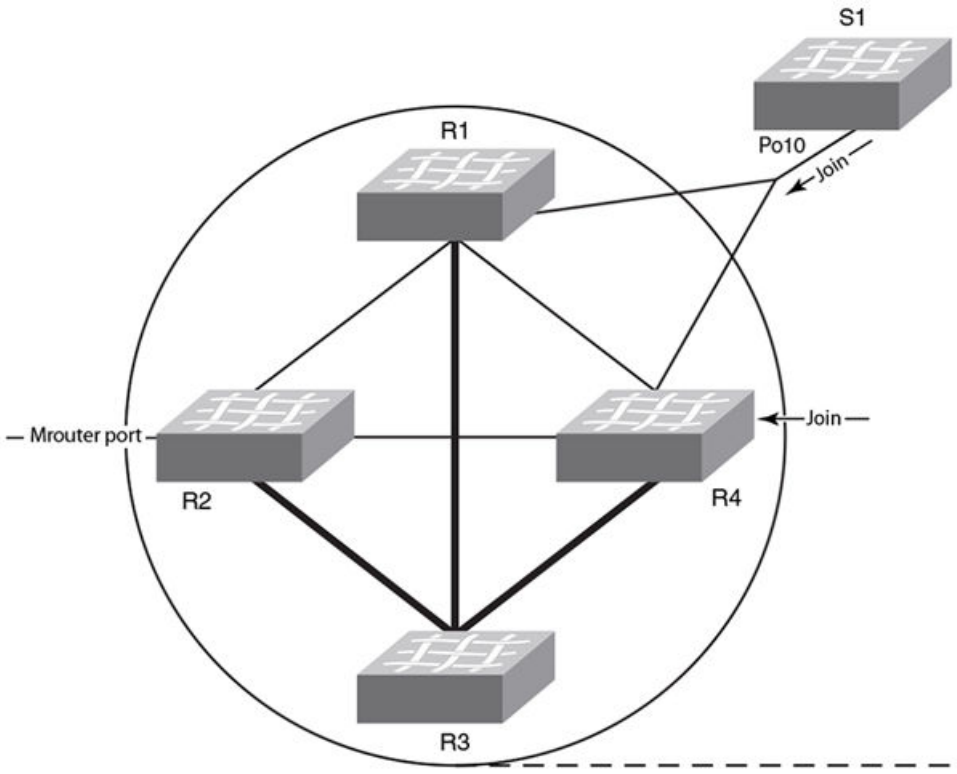

PIM multicast router presence detection

The PIM hello-based multicast router presence detection feature scans the network traffic for incoming PIM hellos.

This feature is enabled by default.

When a PIM hello is detected, that port is marked for the presence of a multicast router and the information is saved. This prevents unnecessary flooding if the PIM designated router (DR) goes offline, as IGMP reports are forwarded to the multicast routers and not only the snooping-enabled router.

IGMP snooping scalability

Here are the scalability limits of IGMP snooping feature in various modes of switch operation for Network OS 4.1.0 and later. The table explains the various metrics involved in describing the scalability limits.

IGMP Metric	Description
Maximum number of IGMP groups supported	This metric is based on the available hardware resources, such as multicast group ID (MGID), configuration replay, and Ethernet Name Server (eNS) distribution bandwidth.
Maximum number of VLANs supported with IGMP snooping configuration	This metric is limited by the general-query packet-generation capacity of IGMP software processes running on the switch, as well as by eNS distribution bandwidth.
Maximum IGMP packet-processing rate per switch	The scalability number described by this metric suggests the upper limit on the number of packets that can be processed by IGMP software processes running on the switch. If the packets are incoming from multiple ports/VLANs, the same processing bandwidth is shared.
Maximum IGMP packet-processing rate per Brocade VCS Fabric cluster	This metric specifies the upper limit on the maximum rate of IGMP packets incoming to a logical Brocade VCS Fabric switch. It is limited by the eNS distribution bandwidth and the number of nodes in the Brocade VCS Fabric cluster.

IGMP snooping in a Brocade VCS Fabric

When supporting a flat Layer 2 network in a data center, VDX switches can be connected in any order to form a cluster. The number of nodes involved in a cluster ranges from four nodes to 24 nodes. Metrics are detailed in the following tables.

TABLE 1 IGMP snooping: four-node cluster metrics

Metric	Limit	Comments
Maximum number of IGMP groups supported	6000	Join requests are sent on four ports of the same switch.
Maximum number of VLANs supported with IGMP configuration	512	
Maximum IGMP packet-processing rate per switch	512 packets/sec	
Maximum IGMP packet-processing rate per Brocade VCS Fabric cluster	512 packets/sec	

TABLE 2 IGMP snooping: 24-node cluster metrics

Metric	Limit	Comments
Maximum number of IGMP groups supported	6000	Join requests are sent on four ports of the same switch.
Maximum number of VLANs supported with IGMP configuration	512	
Maximum IGMP packet-processing rate per switch	512 packets/sec	
Maximum IGMP packet processing rate per Brocade VCS Fabric cluster	512 packets/sec	

TABLE 3 IGMP snooping: Brocade VDX 8770-4 and VDX 8770-8 cluster metrics

Metric	Limit	Comments
Maximum number of IGMP groups supported	6000	Join requests are sent on four ports of the same switch.
Maximum number of VLANs supported with IGMP configuration	512	
Maximum IGMP packet processing rate per switch	512 packet/second	
Maximum IGMP packet processing rate per Brocade VCS Fabric cluster	512 packet/second	

TABLE 4 IGMP snooping: IP multicast metrics

Metric	Limit	Comments
Number of Layer 3 forwarding entries	256	
Number of IGMP snooping forwarding entries	6000	
Number of multicast flows	10000	
PIM interfaces supported	32	
IGMP interfaces supported	32	
IGMP snooping interfaces supported	512	
Learning rate for PIM-SM	32 flows/second	
Learning rate for IGMP snooping	512 groups/second	

IGMP snooping configuration considerations

The following configuration considerations apply to IGMP snooping beginning Network OS release 7.0.0.

- You must remove the IGMP snooping static mrouter configuration from all VLANs before upgrading or downgrading from or to the NOS 6.0.2x release.
- IGMP configuration is not supported on VEs. IGMP configuration is supported only on VLAN, router ports and port-channels.
- IGMP configurations supported under router ports are supported under port-channels.
- IGMP snooping querier and static mrouter can co-exist on a VLAN interface.
- IGMP control packets only with TTL value 1 will be processed.
- Querier and m-router configurations can co-exist.
- You must enable snooping at the global and VLAN levels. Enabling snooping at a global level does not enable snooping on any of the VLANs, however disabling snooping at the global level will disable snooping on all VLANs. This behavior is not supported on clusters where some nodes in the cluster are running NOS firmware prior to 7.0.0 and others in the cluster are running the 7.0.0 firmware.
- When upgrading from versions prior to NOS 7.0.0 to NOS 7.0.0, you must ensure the following post the upgrade procedure:
  1. Snooping is enabled at the global and VLAN level.
  2. Any IGMP configurations under VE should be moved to the VLAN level, as IGMP configuration under VE is not supported.

IGMP snooping upgrade and downgrade considerations

The following table lists the IGMP snooping upgrade and downgrade considerations on VLANs across Network OS 6.0.1 and Network OS 7.x.x.

TABLE 5 IGMP snooping upgrade and downgrade considerations

Network OS 6.0.1	Network OS 7.x.x
Fewer than 512 VLANs are configured with global snooping enabled.	IGMP snooping configuration will be present at the global and VLAN levels after upgrade.
Fewer than 512 snooping-enabled VLANs configured.	Global snooping is automatically enabled after upgrade.

TABLE 5 IGMP snooping upgrade and downgrade considerations (continued)

Network OS 6.0.1	Network OS 7.x.x
Global snooping is not enabled. Only VLAN-level snooping is enabled.	
More than 512 snooping-enabled VLANs, irrespective of whether global snooping is enabled or not.	If the number of snooping-enabled VLANs is more than 512 in Network OS 6.0.x , snooping will be disabled on all VLANs. • If global snooping is enabled and more than 512 VLANs are configured, after the upgrade, you must enable snooping on the VLANs. • If global snooping is not enabled, and if there are more than 512 snooping-enabled VLANs, after upgrading, you must enable VLAN-level and global-level IGMP snooping. This is because only 512 snooping-enabled VLANs are supported across the fabric. If this number was more than 512 in Network OS 6.0.x, snooping on VLANs will be disabled during upgrade.
Downgrading from 7.x.x to 6.0.x	Downgrade is not allowed if IGMP snooping is enabled either at the global or VLAN levels. The following messages display when the downgrade command is executed: **Only global snooping enabled** **Message:** - Downgrade is not allowed because global IGMP snooping is enabled. Please disable global IGMP snooping. **Only VLAN level snooping enabled** **Message:** - Downgrade is not allowed because IGMP snooping is enabled on one or more VLANs. Please disable IGMP snooping on VLANs. **Global and VLAN-level snooping enabled** **Message:** - Downgrade is not allowed because IGMP snooping is enabled globally and on one or more VLANs. Please disable global IGMP snooping and VLAN-level IGMP snooping.

Enabling IGMP snooping

Use the following procedure to enable IGMP snooping on a Data Center Bridging (DCB)/Fibre Channel over Ethernet (FCoE) switch.

You must enable snooping at the global and VLAN levels. Enabling snooping at a global level does not enable snooping on any of the VLANs, however disabling snooping at the global level will disable snooping on all VLANs. This behavior is not supported on clusters where few nodes in the cluster are running NOS firmware prior to 7.0.0 and others in the cluster are running the 7.0.0 firmware.

NOTE
From NOS release 7.0.0, IGMP configurations are supported only on VLAN, LAG, vLAG, and physical ports.

1. Enter the **configure terminal** command to access global configuration mode.

```
device# configure terminal
```

2. Enter the **ip igmp snooping enable** command to enable IGMP snooping at the global level.

```
device(config)# ip igmp snooping enable
```

3. Enter the **interface** command to select the VLAN interface number.

```
device(config)# interface vlan 10
```

4. Optional: Activate the default IGMP snooping querier functionality for the VLAN.

```
device(config-Vlan-10)# ip igmp snooping enable
```

Configuring the IGMP snooping querier

If your multicast traffic is not routed because Protocol-Independent Multicast (PIM) and IGMP are not configured, use the IGMP snooping querier in a VLAN.

The IGMP snooping querier sends out IGMP queries to trigger IGMP responses from devices that are to receive IP multicast traffic. The IGMP snooping querier listens for these responses to map the appropriate forwarding addresses.

Beginning with Network OS 7.0.0, IGMP snooping queries go out with source IP addresses as 0.0.0.0 and in a VCS cluster. The RBridge with the lowest RBridge ID gets elected as the querier.

Use the following procedure to configure the IGMP snooping querier.

1. Enter the **configure terminal** command to access global configuration mode.

   ```
   device# configure terminal
   ```

2. Enter the **interface** command to select the VLAN interface number.

   ```
   sw0(config)# interface vlan 50
   ```

3. Set the IGMP query interval for the VLAN.

   ```
   sw0(config-Vlan-50)# ip igmp query-interval 125
   ```

   The valid range is 1 through 18000 seconds. The default is 125 seconds.

4. Set the last member query count.

   ```
   sw0(config-Vlan-50)# ip igmp last-member-query-count 3
   ```

   The valid range is from 2 through 10. The default is 2.

5. Set the startup query count.

   ```
   sw0(config-Vlan-50)# ip igmp startup-query-count 3
   ```

   The valid range is from 1 through 10. The default is 1.

6. Set the startup query interval.

   ```
   sw0(config-Vlan-50)# ip igmp startup-query-interval 200
   ```

   The valid range is from 1 through 450 seconds. The default is 1 second.

7. Set the Maximum Response Time.

   ```
   sw0(config-Vlan-50)# ip igmp query-max-response-time 10
   ```

   The valid range is from 1 through 25 seconds. The default is 10 seconds.

8. Configure the static Mrouter port.

   ```
   sw0(config-Vlan-50)# ip igmp mrouter interface tengigabitethernet 58/0/7
   ```

9. Configure a static IGMP group.

   ```
   sw0(config-Vlan-50)# ip igmp static-group 225.0.0.1 interface tengigabitethernet 58/0/7 source
   1.0.0.3
   ```

10. Configure the IGMP version.

```
sw0(config-Vlan-50)# ip igmp version v3
```

NOTE
Version 2 is enabled by default. When you change the version of IGMP snooping any existing static or dynamic group will get deleted. These groups will are relearnt at the next query interval when the query is sent out.

11. Set the snooping robustness variable.

```
sw0(config-Vlan-50)# ip igmp robustness-variable 5
```

The valid range is from 2 through 10. The default is 2.

12. You can stop the flooding of the unknown multicast traffic using the **ip igmp snooping restrict-unknown-multicast** command.

```
sw0(config-Vlan-50)# ip igmp restrict-unknown-multicast
```

13. Use the **ip igmp snooping fast-leave** command to enable fast leave processing.

```
sw0(config-Vlan-50)# ip igmp fast-leave
```

14. Activate the IGMP snooping querier functionality for the VLAN.

```
sw0(config-Vlan-50)# ip igmp snooping querier enable
```

NOTE
The IGMP snooping querier and the static mrouter can be configured together on a VLAN interface.

Monitoring IGMP snooping

Monitoring the performance of your IGMP traffic allows you to diagnose any potential issues on your device. This helps you utilize bandwidth more efficiently by setting the device to forward IP multicast traffic only to connected hosts that request multicast traffic.

Use the following commands to monitor IGMP snooping on the device; the commands do not need to be entered in any specific order.

1. Enter the **configure terminal** command to access global configuration mode.
2. Enter the **show ip igmp groups** command to display all information on IGMP multicast groups for the device. Use this command to display the IGMP database, including configured entries for all groups on all interfaces, all groups on specific interfaces, or specific groups on specific interfaces.

```
device# show ip igmp groups
Total Number of Groups: 2
IGMP Connected Group Membership
Group Address  Interface Uptime      Expires      Last Reporter   Version
225.1.1.1      vlan25    00:05:27     00:02:32     25.1.1.1202
Member Ports: eth 2/24
```

3. Use the **show ip igmp statistics** command to display the IGMP statistics for a VLAN or interface.

```
device# show ip igmp snooping statistics interface vlan 10
```

4. Use the **show ip igmp snooping mrouter** command to display mrouter port-related information for all VLANs, or a specific VLAN.

```
device# show ip igmp snooping mrouter
device# show ip igmp snooping mrouter vlan 10
```

When you have reviewed the IGMP statistics for the device, refer to Enabling IGMP snooping on page 17 or Configuring the IGMP snooping querier on page 18 to make any needed corrections.

Using additional IGMP commands

The following commands provide additional support for basic IGMP functionality. For details, refer to the *Brocade Network OS Command Reference*.

Command	Description
ip igmp immediate-leave	Removes a group from the multicast database. Use this command to treat the interface as if it had one multicast client, so that a receipt of a Leave Group request on the interface causes the group to be immediately removed from the multicast database. This command is available for router ports and Layer 3 port channels.
ip igmp snooping fast-leave	Removes a group from the multicast database. This command is available for VLANs.
ip igmp snooping restrict-unknown-multicast	Deactivates or reactivates the flooding of unregistered multicast data traffic on IPv4 IGMP snooping-enabled VLANs. This command functions only with an IPv4 or IPv6 multicast hardware profile.
ipv6 mld snooping restrict-unknown-multicast	Deactivates or reactivates the flooding of unregistered multicast data traffic on IPv6 MLDv1 IGMP snooping-enabled VLANs.
ip igmp static-group	Configures the static group membership entries for a specific interface.

IPv4 Multicast Routing

PIM-sparse overview

PIM-sparse is most effective in large networks sparsely populated with hosts interested in multicast traffic, with most hosts not interested in all multicast data streams.

PIM-sparse devices are organized into domains. A PIM-sparse domain is a contiguous set of devices that all implement PIM and are configured to operate within a common boundary.

PIM-sparse creates unidirectional shared trees that are rooted at a common node in the network called the rendezvous point (RP). The RP acts as the messenger between the source and the interested hosts or routers. There are various ways of identifying an RP within a network. An RP can be configured either statically per PIM router, or by means of a bootstrap router (BSR). Within a network, the RP must always be upstream from the destination hosts.

Once the RP is identified, interested hosts and routers send join messages to the RP for the group in which they are interested. To reduce the number of Join messages incoming to an RP, the local network selects one of its upstream routers as the designated router (DR). All hosts below a DR send IGMP join messages to the DR. The DR sends only one join message to the RP on behalf of all its interested hosts.

PIM-sparse also provides the option of creating a source-based tree rooted at a router adjacent to the tree. This provides the destination hosts with an option of switching from the shared tree to the source-based tree if the latter has a shorter path between the source and the destination.

Bootstrap Router Protocol

For PIM Sparse Mode to function, every PIM router must know the RP in the network, so that it can map multicast groups to the available RP addresses. Bootstrap Router (BSR) Protocol is a mechanism by which a PIM router learns the RP information.

The RP addresses are used as the root of a multicast group-specific distribution tree, the branches of which extend to all the nodes interested in receiving the traffic for that particular multicast group. For multicast sources to reach all receivers, the RP information is crucial so that all PIM routes use the same group-to-RP address mapping. Each node learns the same RP information using the following methods:

- Statically configuring the RP information on each PIM router.
- Using the BSR protocol, which distributes the RP information to each PIM router.

Some of the PIM routers act as Candidate RPs (C-RPs), out of which one C-RP gets elected and acts as RP for a particular group range. In addition, some PIM routers are configured as Candidate BSRs (C-BSRs), and one of these routers will be elected to act as the Bootstrap Router. All PIM routers learn the elected BSR through Bootstrap Messages (BSMs). All Candidate RPs will then report to the elected BSR, which will form the RP-set available in the network and distribute it to all the PIM routers. Therefore all PIM routers eventually have the same RP-set information.

The BSR protocol mechanism converges in the following phases:

- BSR election
- Candidate RP Advertisement and RP set formation
- RP-set distribution

FIGURE 2 BSR Election

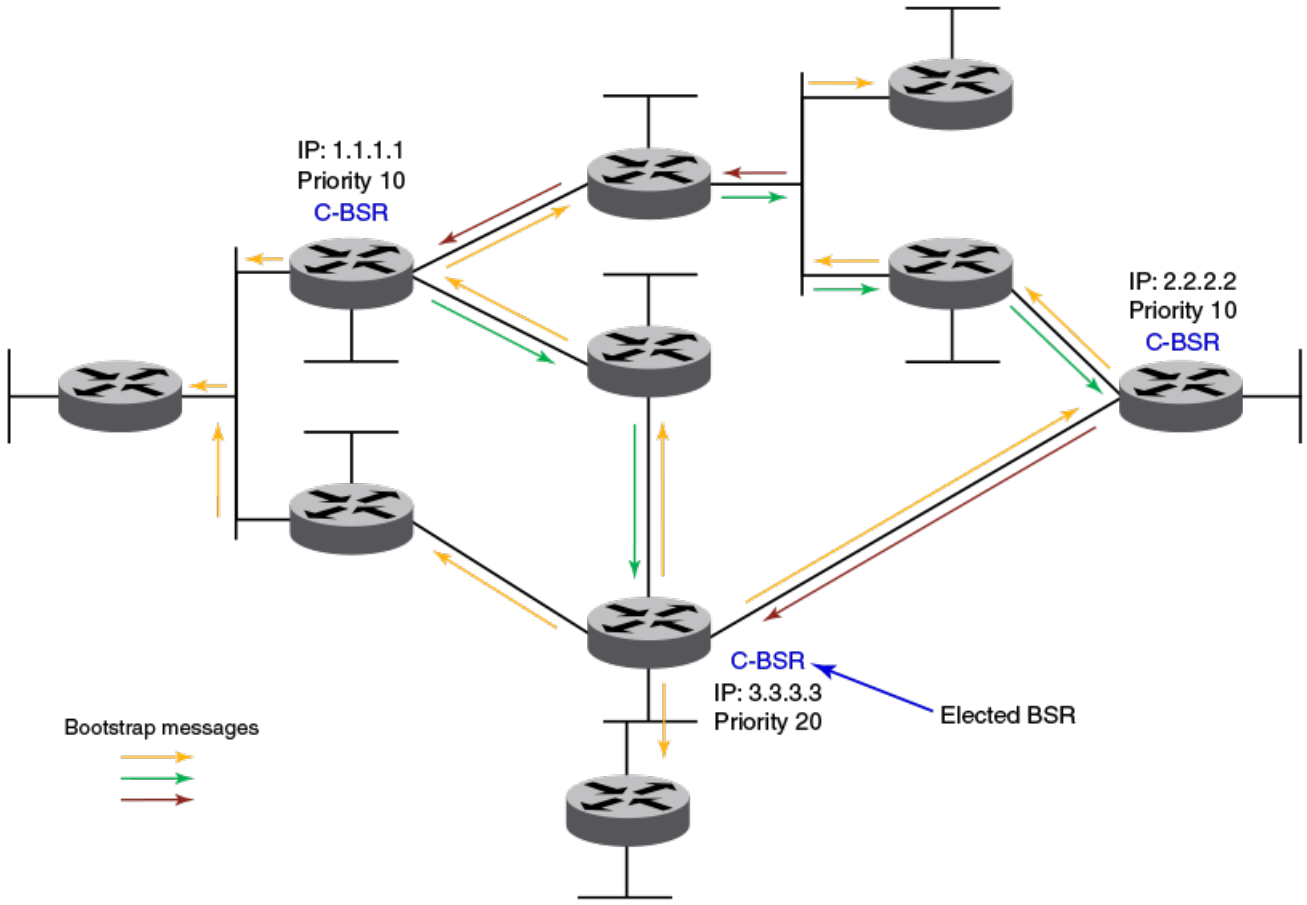

BSR election – Each candidate BSR periodically generates a Bootstrap message (BSM), which carries the configured BSR priority. Every PIM router in the domain floods these BSMs. Other C-BSRs that receive a BSM with higher priority suppress their own BSMs. Eventually, there will be only one C-BSR with BSMs that flood periodically into the network. This single C-BSR becomes the elected Bootstrap Router and its BSM informs all routers that it is the elected BSR.

FIGURE 3 Candidate RP advertisement and RP-set formation

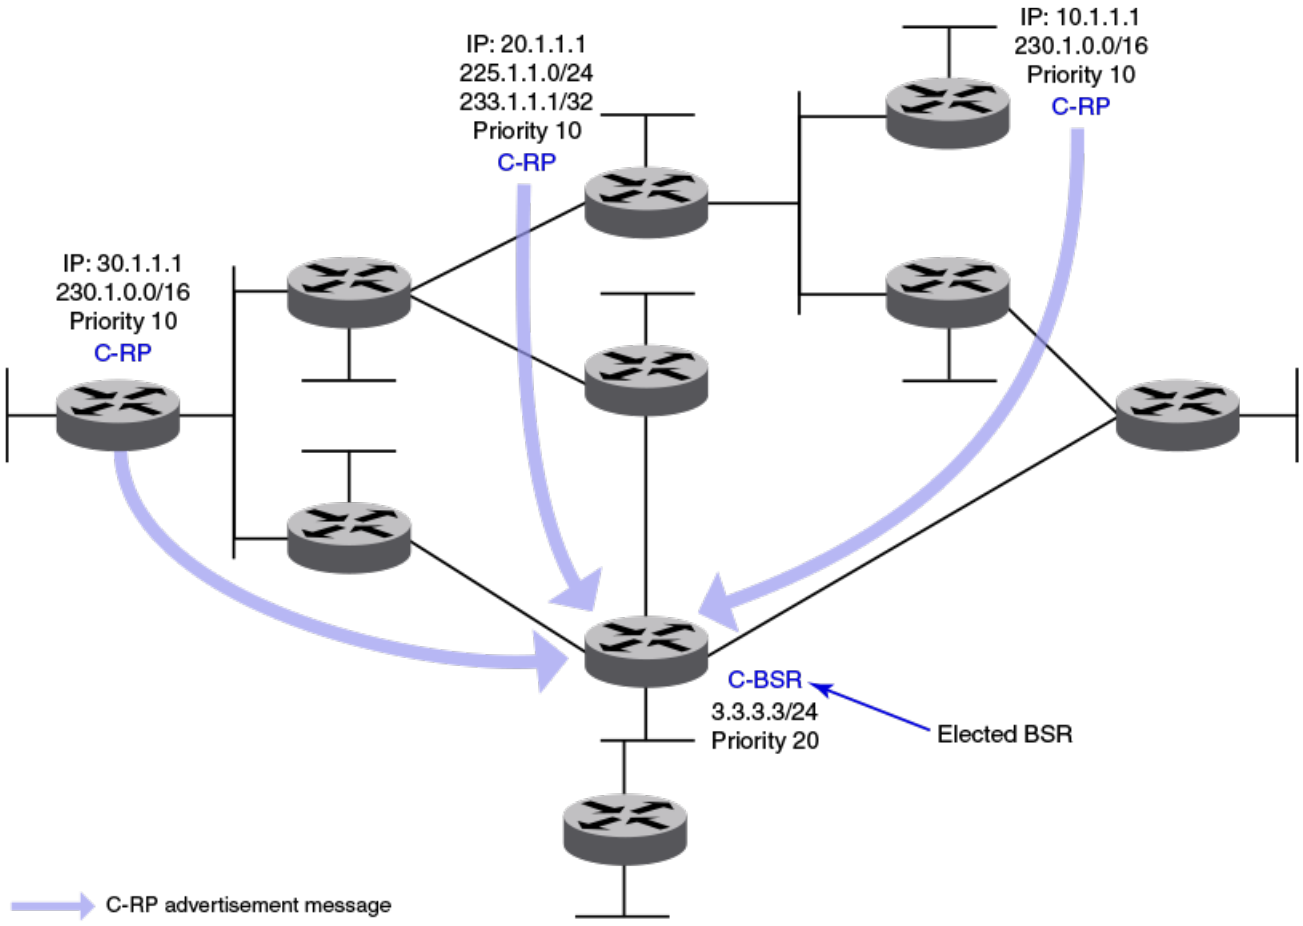

Candidate RP advertisement and RP-set formation: Each candidate RP sends out periodic candidate RP advertisements (C-RP-Adv) messages to the elected BSR. These advertisement messages contain the candidate's priority and a list of multicast group ranges for which this C-RP would like to act as the RP. In addition, it also carries a hold time, for the BSR to discard this C-RP if the hold time expires. In this way, the elected BSR learns about all C-RPs up and reachable. As soon as the BSR starts receiving C-RP advertisements, it builds the RP-set information. This RP-set contains the list for multicast group ranges and C-RP addresses available for each of these group ranges, along with their respective priorities and hold times.

FIGURE 4 RP-set distribution

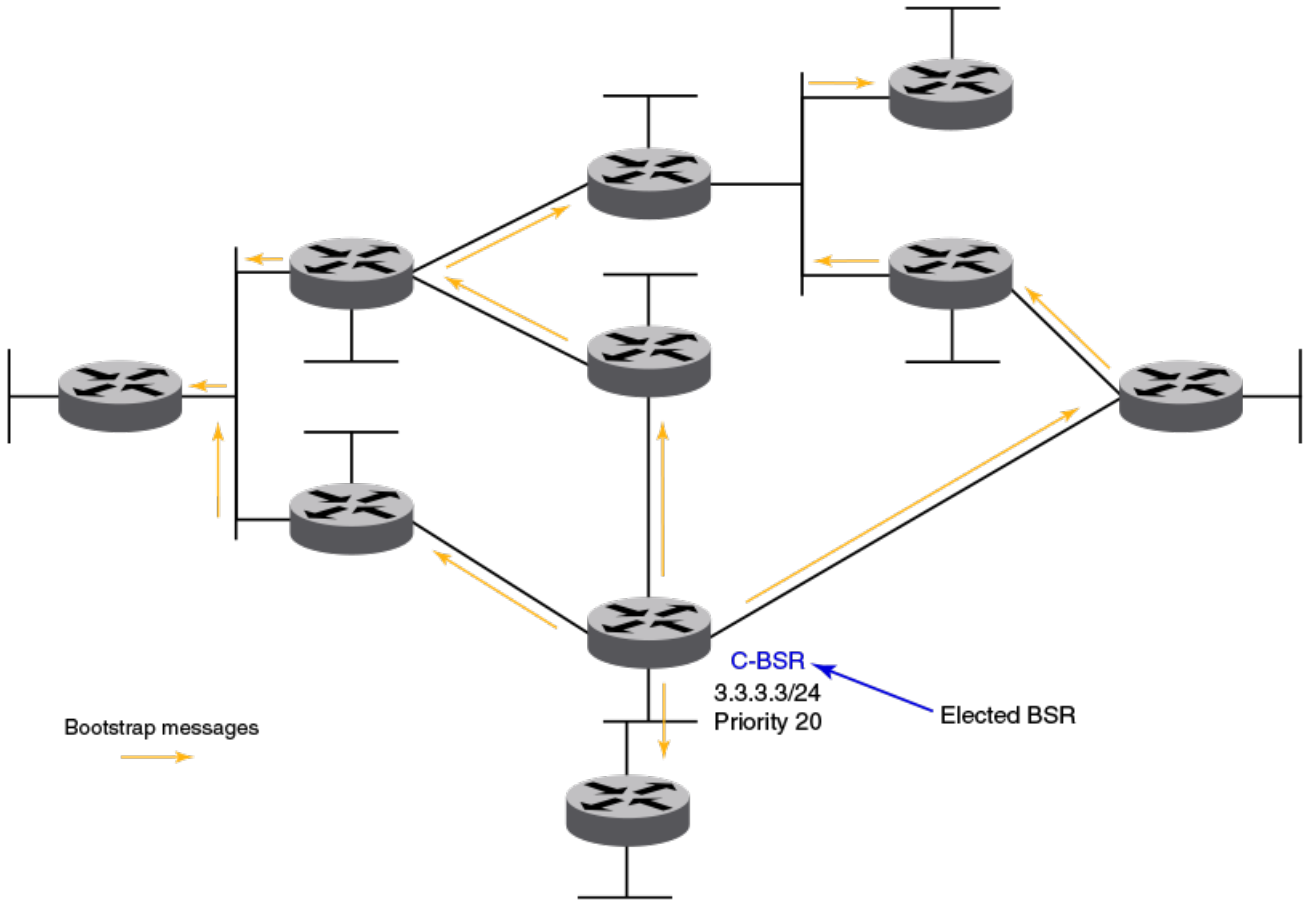

RP-set distribution: The RP-set built by the BSR is set through the same BSM message. Because these BSMs are flooded, the RP-set information rapidly reaches each PIM router. When a PIM router receives the RP-set, it adds all group-to-RP mappings to its pool of mappings, created from static RP configurations as well. Every PIM router runs the same RP hash algorithm to ensure the same C-RP is elected for a particular multicast group throughout the domain. In this way, all PIM routers can build the multicast group-specific distribution tree rooted to the same RP.

BSR timers and values

The BSR mechanism uses timers listed in the following table to ensure the protocol provides reliability and faster convergence. These timers can be configured.

TABLE 6 BSR timers and values

Timer	Default value	Description
Bootstrap message interval	60 seconds	The periodic interval after which a BSM is generated by a BSR.
Bootstrap timeout	130 seconds	The interval after which a BSR is timed out if no BSM is received from it.
Bootstrap minimum interval	10 seconds	The minimum interval after which a BSM should be sent out by a BSR.

TABLE 6 BSR timers and values (continued)

Timer	Default value	Description
C-RP mapping expiry timer	From message	Hold time from C-RP advertisement message. The hold time for C-RP is 2.5 times the RP advertisement interval.
RP mapping expiry timer	From message	Hold time from BSM.
Candidate RP advertisement interval	60 seconds	Periodic interval after which a C-RP generates an advertisement message to the BSR.

RP election algorithm (group-to-RP hashing)

The RP-set information received from the BSR is stored locally and updated by each PIM router periodically upon receiving BSMs. This RP-set contains the list for group prefixes and the corresponding list for C-RP for each group prefix.

The following steps list the RP election procedure for a particular multicast group address:

1. A longest match look-up is performed on all the group prefixes in the RP-set.
2. If more than one C-RP is found by a longest group prefix match, the C-RP with the lowest priority is elected.
3. If more than one C-RP has the same lowest priority, the BSR hash function is used to elect the RP.
4. If the hash functions return the same hash value for more than one C-RP, the highest IP address C-RP is elected.

Using loopback interfaces as an RP

Because loopback interfaces are operationally always up, it is preferable to use them as RPs. Beginning with Network OS 7.1.0, all existing PIM-SM protocol features are also supported on loopback interfaces. Layer 3-enabled loopback interfaces can act as static RP or Candidate-RP. They can also be configured as candidate-BSRs.

PIM-sparse device types

Devices configured with PIM-sparse interfaces also can be configured to fill one or more of the following roles:

- PIM multicast border router (PMBR) — A PIM device that has interfaces within the PIM domain and other interfaces outside the PIM domain. PBMRs connect the PIM domain to the Internet.
- Bootstrap router (BSR): A router that distributes rendezvous point (RP) information to the other PIM-sparse devices within the domain. Each PIM-sparse domain has one active BSR. For redundancy, you can configure ports on multiple devices as candidate BSRs. The PIM-sparse protocol uses an election process to select one of the candidate BSRs as the BSR for the domain. The BSR with the highest BSR priority (a user-configurable parameter) is elected. If the priorities result in a tie, then the candidate BSR interface with the highest IP address is elected.

  The BSR must be configured as part of the Layer 3 core network.
- Rendezvous point (RP): The meeting point for PIM-sparse sources and receivers. A PIM-sparse domain can have multiple RPs, but each PIM-sparse multicast group address can have only one active RP. PIM-sparse devices learn the addresses of RPs and the groups for which they are responsible from messages that the BSR sends to each of the PIM-sparse devices.

  The RP must be configured as part of the Layer 3 core network.

  NOTE
  Brocade recommends that you configure the same ports as candidate BSRs and RPs.

- PIM designated router (DR): Once the RP has been identified, each interested host or router sends join messages to the RP for the group in which they are interested. The local network selects one of its upstream routers as the DR. All hosts below a DR send IGMP join messages to the DR. The DR sends only one join message to the RP on behalf of all its interested hosts. The RP receives the first few packets of the multicast stream, encapsulated in the PIM register message, from the source hosts. These messages are sent as a unicast to the RP. The RP de-encapsulates these packets and forwards them to the respective DRs.

  NOTE
  DR election is based first on the router with the highest configured DR priority for an interface (if DR priority has been configured), and based next on the router with the highest IP address. To configure DR priority, use the **ip pim dr-priority** command.

PIM BSR configuration considerations

The following considerations apply to the BSR protocol implementation in PIM sparse mode:

- The PIM-SM mode must be configured on BSR and RP candidate interfaces.
- PIM neighborship must be maintained for BSR message flooding.
- The unicast route to the BSR candidate and the RP candidate must be resolved.
- Only one BSR candidate and one RP candidate can be configured per node.
- The same interface can act as the BSR and RP candidate.
- The Brocade implementation supports only IPv4 PIM BSR.
- VRFs are not supported.

Protocol-Independent Multicast overview

The Protocol-Independent Multicast (PIM) protocol is a family of IPv4 multicast protocols. PIM does not rely on any particular routing protocol for creating its network topology state. Instead, PIM uses routing information supplied by other traditional routing protocols, such as Open Shortest Path First, Border Gateway Protocol, and Multicast Source Discovery Protocol.

PIM messages are sent encapsulated in an IP packet with the IP protocol field set to 103. Depending on the type of message, the packet is either sent to the PIM All-Router-Multicast address (224.0.0.13) or sent as unicast to a specific host.

As with IP multicast, the main use case of PIM is for the source to be able to send the same information to multiple receivers by using a single stream of traffic. This helps minimize the processing load on the source, as the source needs to maintain only one session irrespective of the number of actual receivers. It also minimizes the load on the IP network, because the packets are sent only on links that lead to an interested receiver.

Several types of PIM exist, but Brocade supports only PIM sparse mode (PIM-sparse, PIM-SM). PIM-sparse explicitly builds unidirectional shared trees rooted at a rendezvous point (RP) per group, and optionally creates shortest-path trees per source.

PIM considerations and limitations

This release includes the following PIM support:

- 8 rendevous point (RP) devices

- 32 virtual interfaces. The virtual interfaces can be either Layer 3 VLAN or router ports
- 32 output interfaces
- 4,000 Layer 3 multicast group IDs
- 2,000 (S,G) forwarding entries
- 256 (*, G) forwarding entries
- A learning rate of 32 routes per second
- Support for high availability (HA)

The following PIM features are not supported:

- Nonstop routing (NSR)
- IP version 6
- VRF

PIM-sparse configuration notes

Be aware of the following issues when configuring PIM-sparse:

- vLAGs must belong to PIM-enabled VLANs. For more information, refer to the "Configuring Link Aggregation" chapter of the *Brocade Network OS Layer 2 Switching Configuration Guide*.
- Set up your vLAGs before performing any PIM-specific configuration.
- Make sure that the rendezvous point (RP) is configured. This functionality should be provided by another router that is outside of the VCS Fabric in Network OS 7.0.1 and prior releases. The RP can be configured either inside or outside the VCS.

  NOTE
  RP functionality, either static or dynamic, is supported for RP inside and outside the VCS Fabric.

- All PIM-enabled aggregation layer devices should have a direct Layer 3 connection to the RP.
- Make sure that the bootstrap router (BSR), if applicable to your setup, is configured. The BSR can be any third-party box that supports PIM, BSR, and rendezvous point (RP) functionality.

PIM-sparse topologies

The figure below shows the components for a single-tier VCS PIM topology.

FIGURE 5 Single-tier VCS deployment

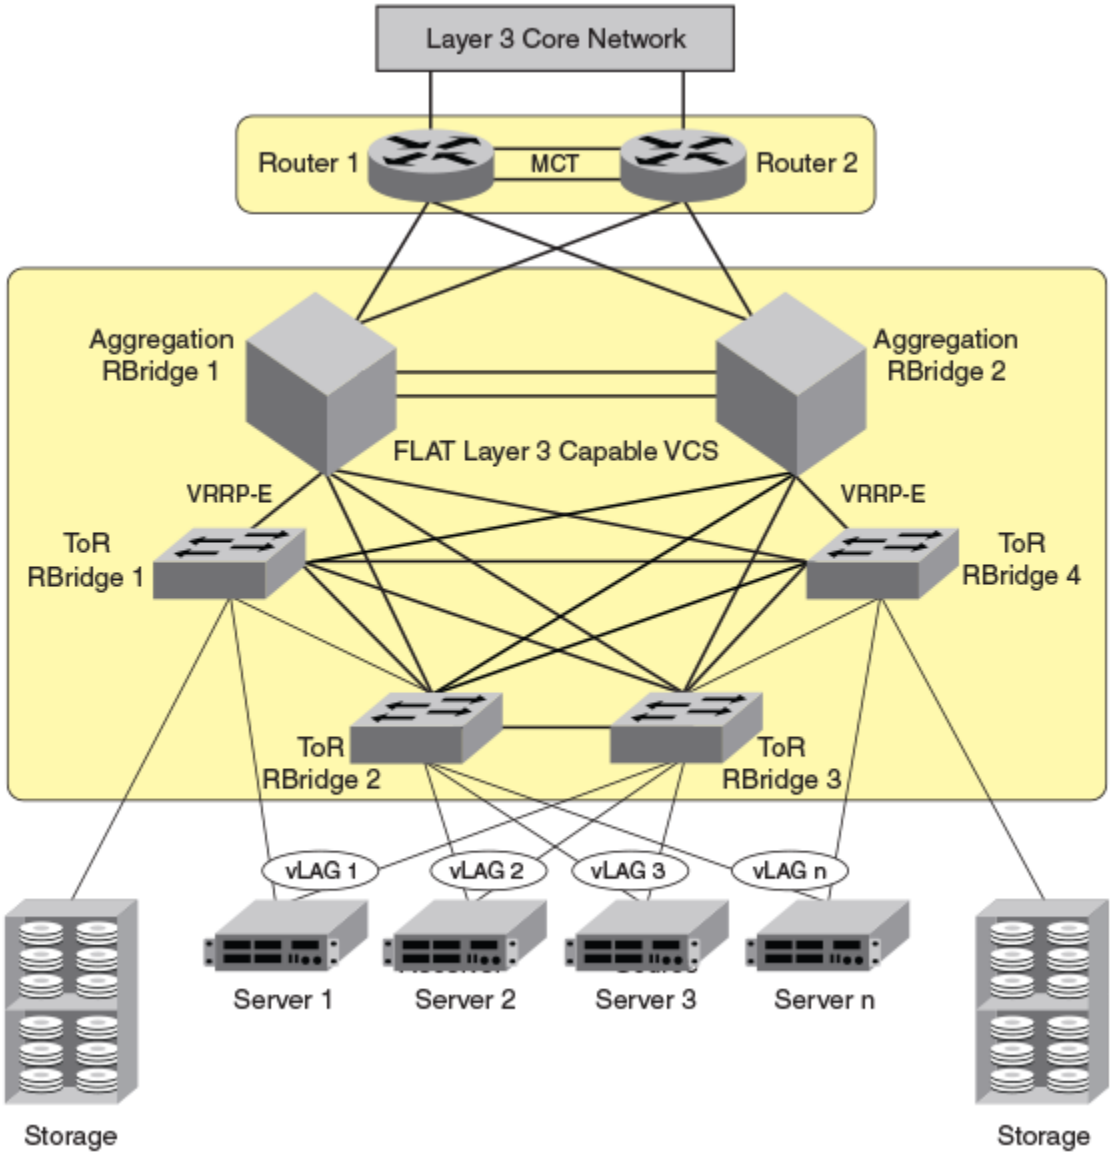

The following requirements apply to the single-VCS deployment depicted in the figure above:

- Top of rack (ToR) switches can be Brocade VDX 6740, VDX 6740T, VDX 6740T-1G, or VDX 8770 models. However, ToR switches are typically only Layer 2-capable when used in this context as part of a PIM environment.
- ToR switches must have IGMP-snooping enabled.
- Layer 3 protocols (VRRP-E and OSPF) can be configured on all interfaces with Layer 3 connectivity to the data-center core.
- IGMP snooping must be enabled on the aggregation-layer switches.
- PIM DR-priority is configured on the virtual Ethernet (VE) interfaces of all PIM-capable aggregation routers to optimize load-sharing abilities within the aggregation layer.

The figure below shows the components for a two-tier VCS PIM topology.

FIGURE 6 Two-tier VCS deployment

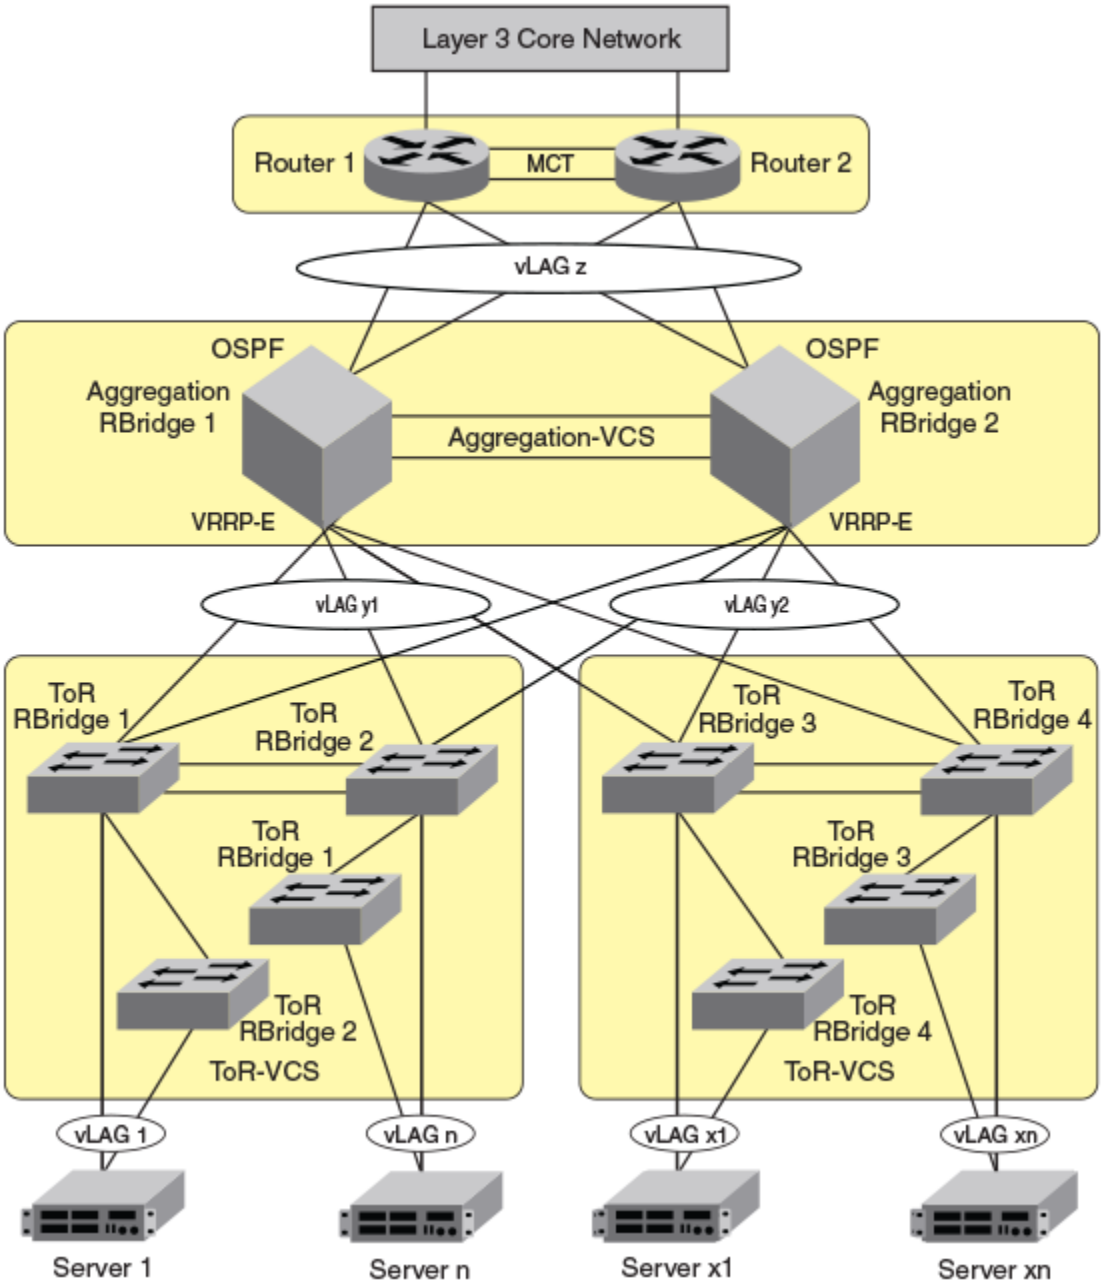

The following requirements apply to the two-tier-VCS deployment depicted in the figure above:

- ToR switches can be Brocade VDX 6740, VDX 8770 or VDX 6940 models. However, ToR switches are typically only Layer 2-capable when used in this context as part of a PIM environment.
- ToR VCS are typically only Layer 2 capable.
- ToR switches must have IGMP snooping enabled.
- Aggregation-layer VCS must be VDX 8770 or VDX 6740 models only.
- Aggregation-layer switches can be PIM-enabled.

- Layer 3 protocols (VRRP-E and OSPF) can be configured on all interfaces with Layer 3 connectivity to the data-center core.
- IGMP snooping must be enabled on the aggregation-layer switches.
- PIM DR-priority is configured on the VE interfaces of all PIM-capable aggregation routers to optimize load-sharing abilities within the aggregation layer.

Graphic guide to PIM-sparse configuration

This diagram represents a sample PIM deployment on a single-tier VCS Fabric.

This network scenario can be used when configuring IGMP snooping on access-layer switches or when enabling PIM on aggregation-layer switches.

FIGURE 7 PIM deployment on a single-tier VCS Fabric example

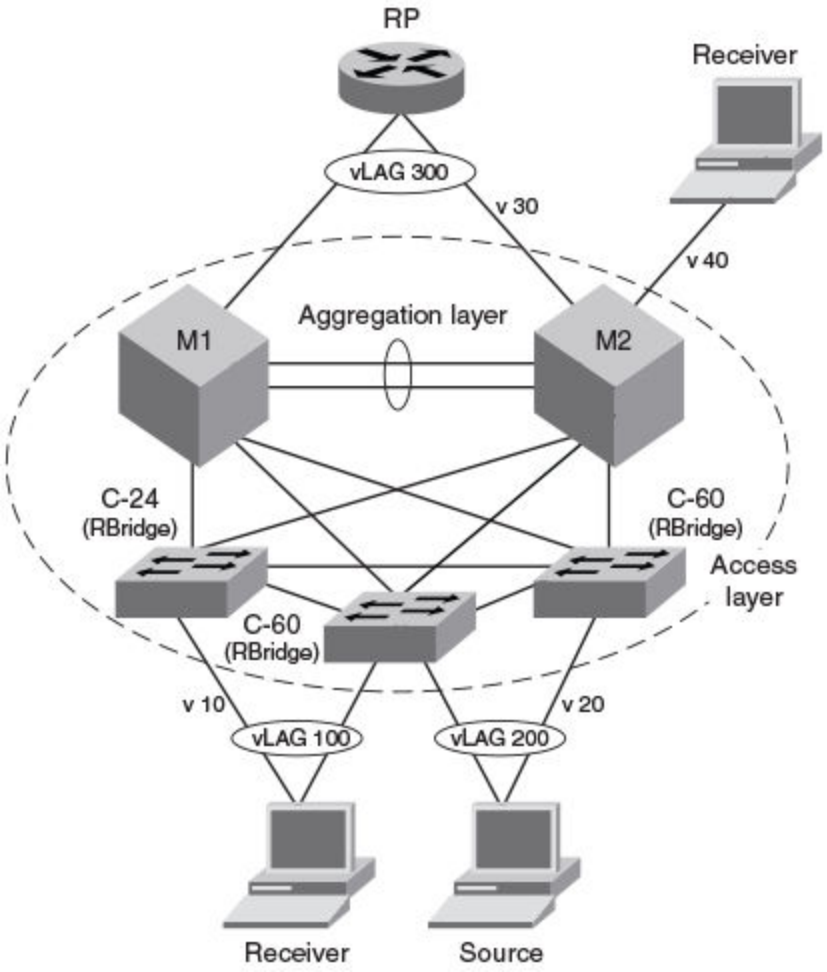

Explanation:

- M1 and M2 can be any variant of Brocade VDX 8770 or VDX 6740 switches.
- M1 is the designated router (DR) for VLAN 10 (labeled "v10") and VLAN 30 (labeled "v30").
- M2 is the designated router (DR) for VLAN 20 (labeled "v20") and VLAN 40 (labeled "v40").

NOTE
Physical interfaces are also supported.

- The switches labeled C-24 and C-60 can be any combination of Brocade VDX 6740 or VDX 8770 models. These switches are pure Layer 2 devices and need IGMP snooping enabled only.

RP within VCS topology

RP within VCS reduces the packet flooding inside the VCS while reaching the RP. Beginning with Network OS 7.1.0, RP within VCS and outside of VCS are supported. With RP within VCS, packets are forwarded onto the exact next hop of the RP. Packets will be forwarded directly onto the physical port when the RP's next hop falls toward the physical port. Packets are forwarded on the port-channel link through which the RP is reachable when the RP's next hop falls towards a port-channel present on the same switch. When the RP's next hop falls towards the VLAG, where a port-channel is present in the TRILL domain, packets are forwarded onto the RBridge (through which the RP is reachable) through the TRILL port.

The following figure shows an example topology representing RP within VCS.

FIGURE 8 RP within VCS topology

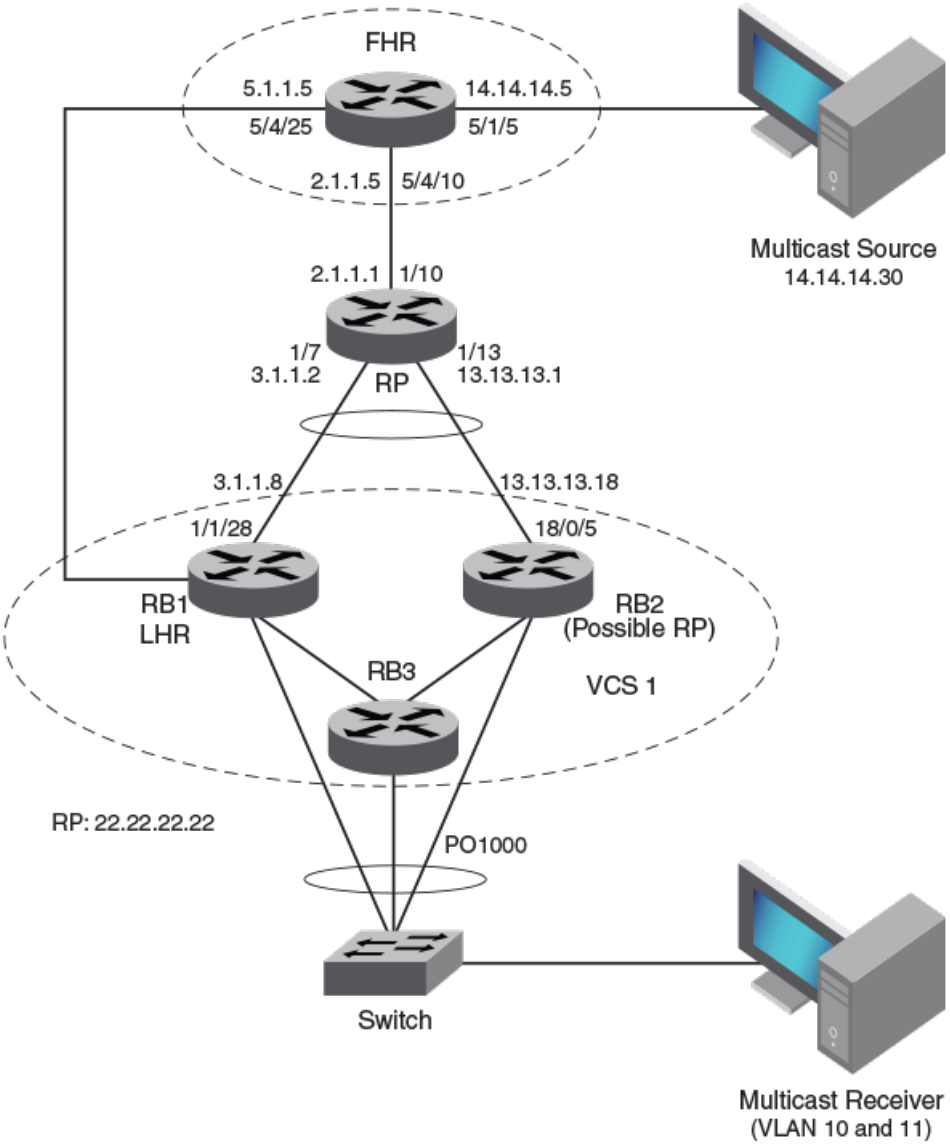

Enabling PIM on a router

Use the following procedure to enable PIM globally.

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter RBridge ID configuration mode.

```
device(config)# rbridge-id 51
device(config-rbridge-id-51)#
```

3. Enter the router PIM configuration mode.

```
device(config-rbridge-id-51)# router pim

device(config-pim-router)#
```

Configuring PIM

Once you enable PIM on a device, you can configure a variety of options in the router PIM configuration mode.

1. Enter the **hello-interval** command to configure the PIM hello timeout.

```
device(config-pim-router)# hello-interval 40
```

2. Enter the **nbr-timeout** command to configure the PIM neighbor timeout.

```
device(config-pim-router)# nbr-timeout 160
```

3. Enter the **bsr-candidate** command to configure the BSR candidate.

```
device(conf-pim-router)# bsr-candidate interface tengigabitethernet 1/1/2 hash-mask-length 32 bsr-
priority 10
```

4. Enter the **bsr-msg-interval** command to configure the BSR message interval.

```
device(config-pim-router)# bsr-msg-interval 30
```

5. Enter the **rp-candidate interface** command to add or remove a candidate RP configuration.

```
device(config-pim-router)# rp-candidate interface ve 101 priority 10
```

   VDX 8770-4 and VDX 8770-8 cannot be used as RP-candidates.

6. Enter the **rp-adv-interval** command to configure the candidate RP advertisement message interval.

```
device(config-pim-router)# rp-adv-interval 30
```

7. Enter the **rp-candidate group-range** command to add or remove a group-range for candidate RP advertisement.

```
device(config-pim-router)# rp-candidate group-range 230.1.0.0 16
```

8. Enter the **message-interval** command to configure the PIM join or prune interval.

```
device(config-pim-router)# message-interval 180
```

9. Enter the **spt-threshold** command to configure the PIM Shortest Path Tree (SPT) threshold.

```
device(config-pim-router)# spt-threshold 1234
```

   NOTE
   In this example the device switches to SPT after receiving 1234 packets. When this value is set to **infinity**, the device never switches to SPT.

10. To configure the IPv4 address of the PIM rendezvous point (RP), use the **rp-address** command.

```
device(config-pim-router)# rp-address 10.22.22.22
```

11. To configure the IPv4 address of the PIM rendezvous point (RP) and also specify a prefix list that defines a multicast group range for the RP.

```
device(config-pim-router)# rp-address 10.22.22.22 prefList1
```

12. To change the inactivity timer, use the **inactivity-timer** command.

```
device(config-pim-router)# inactivity-timer 300
```

13. To disable PIM globally on the device, removing all configurations in PIM router configuration mode:

```
device(config)# rbridge-id 51
device(config-rbridge-id-51)# no router pim
```

Displaying IP PIM information

The following IP PIM show commands can be entered from any level of the CLI.

Use one of the commands to view IP PIM related information. The commands do not need to be entered in the specified order.

1. Display the active C-BSR and information related to local C-RP.

```
device# show ip pim bsr
PIMv2 Bootstrap information
This system is the elected Bootstrap Router (BSR)
BSR address: 207.95.7.1
Uptime: 00:33:52, BSR priority: 5, Hash mask length: 32
Next bootstrap message in 00:00:20
Next Candidate-RP-advertisement in 00:00:10
RP: 207.95.7.1
group prefixes:
224.0.0.0 / 4
Candidate-RP-advertisement period: 60
```

2. Display the current RP-set information learned from elected BSR.

```
device# show ip pim rp-set
Number of group prefixes Learnt from BSR: 2
Group prefix = 224.0.0.0/4 # RPs expected: 1
# RPs received: 1
RP 1: 134.1.1.1 priority=50 age=0 holdtime=150
Group prefix = 225.1.1.0/24 # RPs expected: 1
# RPs received: 1
RP 1: 47.1.1.1 priority=100 age=0 holdtime=150
```

3. Display the current multicast group to RP mappings.

```
device# show ip pim rp-map
Number of group-to-RP mappings: 254
Group address RP address
----------------------------------------
1 226.1.1.1 134.1.1.1
2 226.1.1.2 134.1.1.1
3 226.1.1.3 134.1.1.1
4 226.1.1.4 134.1.1.1
5 226.1.1.5 134.1.1.1
```

Enabling PIM-sparse on routed interfaces

The following procedure enables PIM-sparse and other options on supported interfaces.

This task enables PIM-sparse and other options on any of the following interfaces: TenGigabitEthernet, FortyGigabitEthernet, port-channel, or virtual Ethernet (VE).

You must enable PIM globally before enabling PIM sparse on the interface. You must also enable IGMP snooping at the global and interface level before enabling PIM-sparse on a VE. Refer to enabling IGMP snooping.

1. To enable PIM-sparse on a TenGigabitEthernet interface, enter the global configuration mode.

```
device# configure terminal
```

2. Specify an interface to enable PIM-sparse.

```
device(config)# interface tengigabitethernet 58/0/7
```

3. In the interface configuration mode enter the **ip pim-sparse** command.

```
device(conf-if-te-58/0/7)# ip pim-sparse
```

4. (Optional) To change the designated router (DR) priority from the default, enter the **ip pim dr-priority** command in interface subtype configuration mode and specify a non-default value:

```
device(conf-if-te-58/0/7)# ip pim dr-priority 200
```

5. (Optional) To specify a prefix list as the neighbor filter, enter the **ip pim neighbor-filter** command in interface subtype configuration mode and specify a prefix list, in this example "prefList1":

```
device(conf-if-te-58/0/7)# ip pim neighbor-filter prefList1
```

   NOTE
   By default, directly connected routers under PIM form neighborship with one another. The prefix list denies or permits routers from neighborship. The prefix list must be already configured by means of the ip prefix-list command.

6. To disable PIM-sparse and associated PIM configurations on the interface:

```
device(conf-if-te-58/0/7)# no ip pim pim-sparse
```

Configuring IGMP on routed interfaces

This task enables a variety of options related to IPv4 Internet Group Management Protocol (IGMP) on the supported interfaces.

PIM must be enabled globally on the router before enabling PIM on router ports. Layer 3 IGMP is enabled as soon as the PIM is enabled on the VE or router port. The following IGMP configurations are allowed on router ports, but not allowed on Virtual Ethernet (VE) interfaces, however, the IGMP configurations applied at the VLAN level (Layer 2) will also apply on corresponding VEs.

Enabling and disabling IGMP immediate-leave

IGMP immediate-leave allows a group entry to be removed immediately from the receiver as soon as a Leave Group message is received, as long as the receiver is the only device on the segment that is subscribed to a group. This minimizes the leave latency of group memberships on an interface, as the device does not send group-specific queries. As a result, the group entry is removed from the multicast forwarding table as soon as a group leave message is received.

To enable this feature, enter the **ip igmp immediate-leave** command in interface subtype configuration mode as in the following example.

```
device(conf-if-te-58/0/7)# ip igmp immediate-leave
```

To disable this feature, enter the **no ip igmp immediate-leave** command in interface subtype configuration mode.

```
device(conf-if-te-58/0/7)# ip igmp immediate-leave
```

Configuring the IGMP last-member query interval

You can set the frequency at which IGMP last-member query messages are sent. This is the interval for the response to a query sent after a host leave message is received from the last known active host on the subnet. The group is deleted if no reports are received in this interval. This interval adjusts the speed at which messages are transmitted on the subnet. Smaller values detect the loss of a group member faster.

> NOTE
> If this interval is not configured explicitly, the value is taken from the robustness variable (configured by means of the **ip ipmp robustness-variable** command).

To change the IGMP last-member query interval from the default, enter the **ip igmp last-member-query-interval** command in interface subtype configuration mode as in the following example.

```
device(conf-if-te-58/0/7)# ip igmp last-member-query-interval 1500
```

To restore the IGMP last-member query interval to the default, enter the **no ip igmp last-member-query-interval** command.

```
device(conf-if-te-58/0/7)# no ip igmp last-member-query-interval
```

Configuring the IGMP last-member query count

You can set the number of times that an IGMP query is sent in response to a host leave message. This is the number of times, separated by the last-member query-response interval (configured by the **ip igmp last-member-query-interval** command), that an IGMP query is sent in response to a host leave message from the last known active host on the subnet.

If this interval is not configured explicitly, the value is taken from the robustness variable (configured by means of the **ip ipmp robustness-variable** command).

To change the IGMP last-member query count from the default, enter the **ip igmp last-member-query-count** command in interface subtype configuration mode as in the following example.

```
device(conf-if-te-58/0/7)# ip igmp last-member-query-count 3
```

To restore the IGMP last-member query count to the default, enter the **no ip igmp last-member-query-count** command.

```
device(conf-if-te-58/0/7)# no ip igmp last-member-query-count
```

Configuring the IGMP query interval

You can configure the frequency at which IGMP host query messages are sent. Larger values cause queries to be sent less often.

To change the IGMP query interval from the default, enter the **ip igmp query-interval** command in interface subtype configuration mode as in the following example.

```
device(conf-if-te-58/0/7)# ip igmp query-interval 200
```

To restore the IGMP query interval to the default, enter the **no ip igmp query-interval** command.

```
device(conf-if-te-58/0/7)# no ip igmp query-interval
```

Configuring the IGMP query maximum response time

You can configure the maximum response time for IGMP queries.

To change the IGMP query maximum response time from the default, use the **ip igmp query-max-response-time** command as in the following example.

```
device(conf-if-te-58/0/7)# ip igmp query-max-response-time 9
```

To restore the IGMP query maximum response time to the default, enter the **no ip igmp query-max-response-time** command.

```
device(conf-if-te-58/0/7)# no ip igmp query-max-response-time
```

Configuring the IGMP robustness variable

A robustness value can be configured to compensate for packet loss in congested networks. This value determines the number of general IGMP queries that are sent before a multicast address is aged out for lack of a response. The default is 2.

To change the IGMP robustness variable from the default, use the **ip igmp robustness-variable** command as in the following example.

```
device(conf-if-te-58/0/7)# ip igmp robustness-variable 5
```

To restore the IGMP robustness variable to the default, enter the **no ip igmp robustness-variable** command.

```
device(conf-if-te-58/0/7)# no ip igmp robustness-variable
```

Configuring the IGMP startup query count

The IP IGMP startup query count is the number of queries that are separated by the startup interval. The default is 1.

To change the IGMP startup query count from the default, use the **ip igmp startup-query-count** command as in the following example.

```
device(conf-if-te-58/0/7)# ip igmp startup-query-count 2
```

To restore the IGMP startup query count to the default, enter the **no ip igmp startup-query-count** command.

```
device(conf-if-te-58/0/7)# no ip igmp startup-query-count
```

Configuring the IGMP startup query interval

You can change the interval between the general queries that are sent by the IGMP querier on startup. The default interval is 1.

To change the IGMP startup query interval from the default, use the **ip igmp startup-query-interval** command as in the following example.

```
device(conf-if-te-58/0/7)# ip igmp startup-query-interval 3
```

To restore the IGMP startup query interval to the default, enter the **no ip igmp startup-query-interval** command.

```
device(conf-if-te-58/0/7)# no ip igmp startup-query-interval
```

Configuring the IGMP static group

You can forward traffic statically for a multicast group onto a specified interface, so that the interface behaves as if IGMP joins were received on that interface.

To enable this functionality, use the **ip igmp static-group** command and specify an IPv4 multicast address to be joined, as in the following example.

```
device(conf-if-te-58/0/7)# ip igmp static-group 225.0.0.10
```

To disable this functionality, enter the **no ip igmp static-group** command.

```
device(conf-if-te-58/0/7)# no ip igmp static-group
```

Restricting unknown multicast

The restrict-unknown-multicast feature prevents the default flooding of multicast traffic on all ports of a VLAN.

The PIM topology and VLANs must be configured before this feature can be activated.

When this feature is enabled, (*,G,V) entries are programmed, and the non-PIM-DR does not process or create (*,*,V) routes and maintain them in the mrouter database. IP multicast data traffic is sent only to mrouter-learned ports or PIM-hello-learned ports.

1. Enter interface configuration mode for the VLAN whose unknown multicast traffic is to be restricted.

   ```
   device(config)# interface vlan 100
   ```

2. Enter the **ip igmp snooping restrict-unknown-multicast** command.

   ```
   device(config-Vlan-100)# ip igmp snooping restrict-unknown-multicast
   ```

Multicast on bridge domain

Bridge domain interface is a logical interface that allows bidirectional flow of traffic.

Multiple service end points like port-vlan or port-vlan-vlan can be made part of a single broadcast domain where we can achieve any-to-any bridging called Bridge Domain. The service end points can be of different types like Pseudo wire, VxLAN tunnel/VNI endpoints and other. Multicast IGMP snooping and IP Multicast PIM are supported on Bridge Domain.

IGMP on bridge domain

IGMP snooping on bridge domain is used to learn the particular multicast group on specific ports associated with the bridge domain by trapping the IGMP control packets and programing the hardware entries with learned Multicast groups and list of interested ports that are part of the bridge domain. When multicast traffic comes from a source the traffic is sent to the interested receivers instead of flooding the multicast traffic on all ports of the bridge domain. IGMP on bridge domain internally works same as it works on VLAN. Bridge domain contains logical interfaces (LIFs) so corresponding multicast groups contain the LIFs as the outgoing interfaces.

PIM support on VE bind to bridge domain

Layer3 interface VE can be bound to VLAN or the bridge domain. PIM on bridge domain works similar to VLAN case where the VE is bound to VLAN or bridge domain. PIM-SM can be configured on the VE. PIM snooping is also supported on Bridge domain. PIM snooping is used to send the data traffic to only interested PIM routers rather than all PIM routers connected in the Bridge domain.

To enable PIM snooping on bridge domain, use enable command.

```
device(config)# bridge-domain 10
device(config-bridge-domain-10)# ip pim-sparse
```

Configuring IGMP snooping on a bridge domain

Follow these steps to configure IGMP snooping on a bridge domain.

1. To enable IGMP snooping on bridge domain, use enable command. To disable the function, use the no form of this command.

   ```
   device # ip igmp snooping enable
   device(config)# bridge-domain 10
   device(config-bridge-domain-10)# ip igmp snooping enable
   ```

   Syntax: [no] ip igmp snooping enable

2. To enable IGMP querier on bridge domain, use enable command. To disable the function, use the no form of this command.

   ```
   device(config)# bridge-domain 10
   device(config-bridge-domain-10)# ip igmp snooping querier enable
   ```

   Syntax: [no] ip igmp snooping querier enable

3. IGMP version.

   ```
   device(config)# bridge-domain 10
   device(config-bridge-domain-10)# ip igmp version v2
   ```

   Syntax: [no] ip igmp snooping version <v1/v2/v3>

4. To enable IGMP fast leave on bridge domain, use enable command. To disable the function, use the no form of this command.

```
device(config)# bridge-domain 10
device(config-bridge-domain-10)# ip igmp snooping fast-leave
```

   Syntax: [no] ip igmp snooping fast-leave

5. Querier interval. The default is 125 seconds.

```
device(config)# bridge-domain 10
device(config-bridge-domain-10)# ip igmp snooping query interval 30
```

   Syntax: [no] ip igmp snooping query interval <1-18000>

6. Query max response time.

```
device(config)# bridge-domain 10
device(config-bridge-domain-10)# ip igmp snooping query-max-response-time 20
```

   Syntax: [no] ip igmp snooping query-max-response-time <1-25>

7. IGMP last membership query interval.

```
device(config)# bridge-domain 10
device(config-bridge-domain-10)# [no] ip igmp snooping last-member-query-interval  150
```

   Syntax:[no] ip igmp snooping last-member-query-interval <100-25500>

The following example is the steps in the previous configuration:

```
device(config-BD-id)# ip igmp snooping ?
Possible completions:
  enable                        IGMP Enable
  fast-leave                    Fast Leave Processing
  last-member-query-interval    Last Member Query Interval
  mrouter                       Multicast Router
  querier                       Querier
  query-interval                Query Interval
  query-max-response-time       IGMP Max Query Response Time
  static-group                  Static Group to be Joined
  version                       IGMP Snooping Version
```

Multi-Chasis Trunk (MCT)

A Multi-Chassis Trunk (MCT) is a trunk that initiates at a single MCT-unaware server or switch and terminates at two MCT-aware switches.

Link Aggregation (LAG) trunks provide link level redundancy and increased capacity. However, LAG trunks do not provide switch-level redundancy. If the switch to which the LAG trunk is attached fails, the entire LAG trunk loses network connectivity. With MCT, member links of the LAG are connected to two chassis. The MCT switches may be directly connected using an Inter-Chassis Link (ICL) to enable data flow and control messages between them. In this model, if one MCT switch fails, a data path will remain through the other switch.

In an MCT scenario, all links are active and can be load shared to increase bandwidth. In addition, traffic restoration can be achieved in milliseconds after an MCT link failure or MCT switch failure. MCT is designed to increase network resilience and performance.

MP-BGP EVPN

Multi-protocol BGP is an extension to BGP that enables BGP to carry routing information for multiple network layers. Ethernet VPN EVPN) connect a group of customer sites using a virtual bridge. Treats MAC addresses as routable addresses and distribute them in Border Gateway Protocol (BGP). Uses Multi-protocol BGP (MP-BGP)

Advantages of MP-BGP EVPN

The advantages of MP-BGP EVPN are -

- It is standard based.
- Scalable and reliable because of Border Gateway Protocol base.
- Policies can be applied.
- Support exchange of IP addresses and IP prefixes.

Layer 2 Multicast Snooping over MCT

Multicast control packets behavior on Multi-Chasis Trunk (MCT).

Multicast state information is synced between the MCT peers using MP-BGP EVPN transport. Multicast protocol packets will not be sent on the peer link unless required.

IGMP/MLD protocol packets are of two types:

1. Membership query
    - General query - In a query message, the multicast address field is set to 0 when MLD sends a general query. The general query learns which multicast addresses have listeners on an attached link.
    - Group specific query - A group address is a multicast address.
2. Membership reports - In a report, the multicast address field is that of the specific multicast address to which the sender is listening.
    - Version 1 membership report - MLD version 1 is based on version 2 of the Internet Group Management Protocol (IGMP).
    - Version 2 membership report - MLD version 2 is based on version 3 of the IGMP. MLD version 2 is fully backward-compatible with MLD version 1.
    - Leave group

IGMP/MLD query packet processing on MCT

For IGMP queries, each EVI has BUM suppressed MGID associated, IGMP/MLD query packets need to be transmitted on ICL to address the following scenarios:

- The querier connected to only one of the MCT peer switch becomes elected querier.
- Only one of the peer switch is configured as querier.
- The switch would age out IGMP/MLD routes if memberships are not confirmed within time out interval.Although query packet is received on MCT peer link, mrouter port is not learnt/considered on that peer link.

IGMP/MLD membership reports

For IGMP reports and leaves,

- Traditionally, each peer switch learns about L2 Multicast memberships by snooping the IGMP/MLD membership reports. The membership reports are then flooded on mrouter ports.
- For MCT, since mrouter port is not learnt on the peer link, membership reports are not flooded between the peer switches. Peer switches exchange learnt routes using EVPN NLRI messages between BGP peers running on MCT cluster control vlan.
- BGP on the peer switch would communicate with the IGMP/MLD routes to Multicast module to add/delete group memberships and allocates/rejects the MGIDs and installs/un-installs routes in hardware.
- Also, each peer switch would generate proxy report for the IGMP/MLD routes learnt from EVPN NLRI, if general query or group-specific query is received from any port, other than peer-link.

Leave membership report

When fast-leave is not configured and MCT peer receives leave membership report from one of its clients for group G, the switch/router informs other MCT peers about the group specific query and latency using IGMP leave synch route. The peer switch, which runs IGMP querier sends group specific query and group query.

mRouter synchronization

mRouter synchronization helps in achieving optimal path selection for unknown multicast traffic and optimal MP-BGP message exchange between MCT Peers.

BGP handling of EVPN IGMP routes

In DC applications, EVPN is used as way of standard inter-POD communication for both intra-DC and inter-DC.

A subnet can span across multiple PODs and DCs. EVPN provides robust multi-tenant solution with extensive multi-homing capabilities to stretch a subnet (e.g., VLAN) across multiple PODs and DCs. There can be many hosts/VMs (e.g., several hundreds) attached to a subnet that is stretched across several PODs and DCs. These hosts/virtual machines express their interests in multicast groups on a given subnet/VLAN by sending IGMP membership reports (Joins) for their interested multicast group(s). Also, an IGMP router (e.g., IGMPv1) periodically sends membership queries to find out if there are hosts on that subnet still interested in receiving multicast traffic for that group. Just like ARP/ND suppression mechanism in EVPN to reduce the flooding of ARP messages over EVPN, it is also desired to have a mechanism to reduce the flood of IGMP messages (both Queries and Reports) in EVPN. This is achieved through IGMP Join Sync and IGMP Leave Sync EVPN routes specified in the draft "draft-ietf-bess-evpn-igmp-mld-proxy-00".

TABLE 7 EVPN Routes

EVPN route type	Route type name	Usage
7	IGMP Join Synch Route	To exchange (S,G)/(*,G) learnt on BGP EVPN peers.
8	IGMP Leave Synch Route	To exchange group leave between BGP EVPN peers.

IGMP Join Synch Route

IGMP allows a network host to inform a router that it is interested in receiving a particular multicast stream.

To begin, the multicast group is assigned a multicast address (that is, an IP address in the 224.0.0.0/4 class D address space). Hosts register to receive the stream join the group by sending an IGMP Report to the upstream multicast router. The router then adds that group to the list of multicast groups that should be forwarded onto the local subnet.

The router does not maintain state about which hosts on the subnet are to receive traffic for the group. Instead, the router continues to send traffic to the subnet until either a timeout value expires or there are no more hosts in that group on the subnet.

TABLE 8 IGMP Join Synch Route

Packets	Usage
RD (8 Octets)	
Ethernet Segment Identifier (10 Octets)	0 if (S, G) / (*, G) is learnt on CEP
Ethernet Tag ID (4 octets)	0 or optionally, set to Vlan ID over which Multicast Route is learnt
Multicast Source Length (1 octet)	32 for IPv4 address, 128 for IPv6 address and 0 for Wildcard address (*)
Multicast Source Address (variable)	IP address of multicast source
Multicast Group Length (1 octet)	32 for IPv4 address, 128 for IPv6 address
Multicast Group Address (variable)	IP address of multicast group
Originator Router Length (1 octet)	32 for IPv4 address, 128 for IPv6 address
Originator Router Address (variable)	The IP address of the originating router.
Flags (1 octets) (optional)	The flag fields are defined below in the table

NOTE

For,
Querier Config Sync - Multicast Group is set to 224.0.0.2.
Mrouter Sync - Multicast Group is set to 224.0.0.1.

TABLE 9 Flags

0	1	2	3	4	5	6	7
reserved	Querier config synch	mRouter synch	PIM snooping	IE	V3	V2	V1
	Bit 1 indicates QuerierConfigSync.	Bit 2 indicates MrouterSync.	Bit 3 indicates PIM Snooping.	Bit 4 indicates whether the (S, G) information carries within the route-type of Include Group type (bit value 0) or an Exclude Group type (bit value 1). The Exclude Group type bit should be ignored if bit 5 is set in case of IGMP Version 2 & IGMP Version 1 & MLD Version 1.	Bit 5 indicates support for IGMP/MLD version 3.	Bit 6 indicates support for IGMP/MLD version 2.	Bit 7 indicates support for IGMP/MLD version 1.

NOTE

Bits 1, 2, 3 are proprietary fields.

IGMP Leave Synch Route

When a host no longer wants to receive multicast traffic, it sends the router an IGMP Leave message.

After receiving this message, the router sends a query to the local subnet to determine whether any group members remain, sending the message to all hosts on the subnet, at the multicast All-Hosts address (224.0.0.1). If any host responds, the router continues to send to the group; if not, the router removes the multicast group from its forwarding list and stops sending to the group.

TABLE 10 IGMP Join Synch Route

Packets	Usage
RD (8 Octets)	
Ethernet Segment Identifier (10 Octets)	0 if (S, G) / (*, G) is learnt on CEP
Ethernet Tag ID (4 octets)	0 or optionally, set to Vlan ID over which Multicast Route is learnt
Multicast Source Length (1 octet)	32 for IPv4 address, 128 for IPv6 address and 0 for Wildcard address (*)
Multicast Source Address (variable)	IP address of multicast source
Multicast Group Length (1 octet)	32 for IPv4 address, 128 for IPv6 address
Multicast Group Address (variable)	IP address of multicast group
Originator Router Length (1 octet)	32 for IPv4 address, 128 for IPv6 address
Originator Router Address (variable)	The IP address of the originating router.
Flags (1 octets) (optional)	The flag fields are defined below in the table

TABLE 11 Flags

0	1	2	3	4	5	6	7
reserved	reserved	reserved	reserved	IE	V3	V2	V1
				Bit 4 indicates whether the (S, G) information carries within the route-type of Include Group type (bit value 0) or an Exclude Group type (bit value 1). The Exclude Group type bit should be ignored if bit 5 is set in case of IGMP Version 2 & IGMP Version 1 & MLD Version 1.	Bit 5 indicates support for IGMP/MLD version 3.	Bit 6 indicates support for IGMP/MLD version 2.	Bit 7 indicates support for IGMP/MLD version 1.

NOTE

Leave Group Synchronization & Maximum Response time are used during in Leave group synch procedures.

IGMP Join and Leave a Multicast Group

IGMP routes are originated by MLD and are sent to BGP through RibLib to be transported to the BGP EVPN Peers.

These routes are carried in BGP EVPN NLRI as type 7 and type 8 routes. BGP adds the EVI-RT extended community to the EVPN NLRI and transports the route to the EVPN peers. EVPN configuration must be present for the specified EVI.

EVPN IGMP routes received from the remote peers are validated against import rules and added to VPN table in BGP only, if the validation passes. The routes are later imported into the MACVRF table, if the RT in the route matches the RTs configured for given EVI (VLAN/BD). Consolidation of the routes from different sources (RDs) happen in BGP after the route passes the route target and import filtering checks.

The routes are then installed in BGP and also forwarded to MLD through RibLib. Installed routes are also forwarded to other EVPN peers. BGP EVPN should be configured to support the IGMP routes.

NOTE
No new configuration is needed in BGP to support the IGMP routes.

EVI Route Target extended community

The EVI Route Target (RT) is a new EVPN extended community of type 6 and sub-type yet to be defined by IANA. However, EVI-RT extended community is NOT supported for IGMP routes.

Instead, the Route Target extended community with Type 0x00 and Sub-type 0x02 is supported. This extended community carries the RT associated with the EVI (VLAN/BD), so that the receiving PE can identify the EVI properly.

ES-Import Route Target extended community

ES-Import (ES-I) Route Target (RT) is another EVPN extended community of Type 0x06 and Sub-type 0x02, the 6 byte value calculated is based on the ES-Import.

The IGMP Join and Leave Synch routes carry the ES-Import RT for the ES on which the IGMP membership report was received. Thus, it may only go to the PEs attached to that ES (and not to any other PEs).

Encap Type support

Only MPLS and NSH tunnel encapsulation types are supported for the IGMP Sync routes. VxLAN tunnel encapsulation type is not supported. Hence, the IGMP Sync routes are not advertised to VxLAN peers.

Traffic Forwarding Path for L2 Multicast

Both peers are updated with (*, G) membership for CCEP.

When a receiver connected on CCEP sends membership report to join group G. The DF election of the CCEP and MGID updation by multicast module prevents duplication of multicast data packets destined to group G on CCEP as well as peer-link.

Always honor DF election while programming receivers on CCEP in MGID. However, the path given by DF election may not be optimal as it might direct the multicast data traffic originating at one peer switch to receiver on CCEP via peer-link even though (*,G) membership does not include CEP on the peer-switch that does not host the source.

When a group is learnt on member VLANs, the DF for IGMP/MLD route is elected by hashing on IVID of the member VLAN, Source-IP and Group-IP.

The OIF list of IGMP/MLD route on DF includes:

- Receivers connected via CCEP.
- ICL if its MCT peer has receivers connected via CEP.
- Receivers connected via local CEP.

The OIF list of IGMP/MLD route on non-DF includes:

- ICL to redirect the multicast stream to DF of the stream.
- Receivers connected via CEP.

Data Encapsulation of L2 Multicast Traffic on ICL

Data Encapsulation of L2 Multicast from CEP/CCEP received on member vlan is similar to L2 Flooding traffic.

- If the CCP is down, it will forward locally.
- If the remote CCEP is down, it will forward locally.
- If the local CCEP is down, it will not forward locally.
- If the ingress is the CEP, it will forward locally.
- If the ingress is the ICL, it will not forward locally.
- If the ingress is a different CCEP, it will forward locally.

Layer 3 multicast on multiple subnets

The Layer 3 multicast on multiple subnets functionality enables Layer 3 PIM to function on all networks or subnets configured on an interface with primary and multiple secondary IP addresses.

In releases prior to Network OS 7.1.0, IP Multicast was supported on Layer 3 interfaces belonging to one network only. If the interface was configured with multiple IP addresses belonging to a different network or subnet, Layer 3 multicast would only function on the primary address. As a result, multicast source or IGMP hosts attached to non-primary networks of any interface were not supported.

With the Layer 3 multicast on multiple subnets functionality, the lowest IP address on an interface will always act as a primary address, and this primary address is used to send out PIM hello messages. Neighborship is built on the primary address only. However, each hello message will carry a list of secondary IP addresses, so that each PIM neighbor is aware of these IP addresses which can be used as an upstream neighbor address for an alternate RPF path.

High availability is supported for the PIM multinet functionality. The multinet enable or disable configuration per interface will be synced to the standby.

FIGURE 9 PIM neighborship

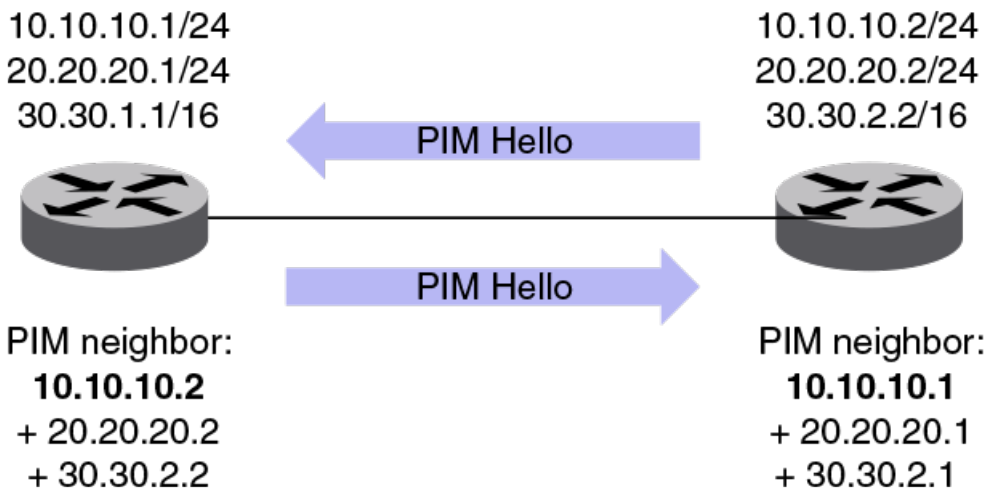

Configuring PIM multinet

PIM multinet can be configured at the interface level.

1. Enter the global configuration mode.

```
device# configure terminal
```

2. Enter the interface subtype configuration mode.

```
device(config)# interface te 1/2/8
```

3. Enter the **ip pim multinet enable** command.

```
device(conf-if-te-1/2/8)# ip pim multinet enable
```

The following example configures PIM multinet on a TenGigabitEthernet interface.

```
device# configure terminal
device(config)# interface te 1/2/8
device(conf-if-te-1/2/8)# ip pim multinet enable
```

The following example displays the PIM neighbors.

```
device# show ip pim neighbor
Total Number of Neighbors : 3
Port         Phy_Port     Neighbor        Holdtime Age         UpTime      Priority   Tracking-Bit
                                          sec      sec      Dd HH:MM:SS
Ve 16        Ve 16        16.1.1.1        105      11       4d 23:55:07         1        Enabled
Ve 26        Ve 26        26.1.1.1        105      25          03:49:29         1        Enabled
                          + 26.1.2.1
                          + 26.1.3.1
                          + 26.1.4.1
                          + 26.1.5.1
Ve 65        Ve 65        65.1.2.2        105      18       5d 22:04:44         1        Enabled
                          + 65.1.3.2
                          + 65.1.4.2
                          + 65.1.5.2
```

IPv6 Multicast VLAN Traffic Reduction

MLD snooping overview

Multicast Listener Discovery (MLD) snooping is a multicast-constraining mechanism that runs on Layer 2 or Layer 3 devices to manage and control IPv6 multicast groups.

A Layer 2 switch forwards all multicast control packets and data received on all the member ports of a VLAN interface. This approach, though simple, is not bandwidth efficient, because only a subset of member ports may be connected to devices interested in receiving those multicast packets. In the worst-case scenario the data are forwarded to all port members of a VLAN with a large number of member ports, even if only a single VLAN member is interested in receiving the data. Such scenarios can lead to loss of throughput for a switch when it receives high-rate multicast data traffic.

MLD snooping provides functionality for IPv6 that is similar to IGMP snooping for IPv4, by sending IPv6 multicast traffic only to interested listeners. By listening to and analyzing MLD messages, a Layer 2 device running MLD snooping establishes mappings between ports and multicast MAC addresses or multicast IP addresses, and forwards multicast data accordingly. Multicast routers in a network are found by means of either static configuration, dynamic learning, or PIM hello-based mrouter detection.

NOTE
This release supports the IPv6 version of MLDv1 snooping.

In any given subnet, one multicast router is elected to act as an MLD querier. The MLD querier sends out the following types of queries to hosts:

- **General query**: Querier asks whether any host is listening to any group.
- **Group-specific query**: Querier asks whether any host is listening to a specific multicast group. This query is sent in response to a host leaving the multicast group and allows the router to determine quickly whether any remaining hosts are interested in the group.

Hosts that are multicast listeners send the following kinds of messages:

- **Report message**: Indicates that the host wants to join a particular multicast group.
- **Done message**: Indicates that the host wants to leave a particular multicast group.

MLD traffic is forwarded as follows:

- MLD general queries received on a multicast-router interface are forwarded to all other interfaces in the VLAN.
- MLD group-specific queries received on a multicast-router interface are forwarded to only those interfaces in the VLAN that are members of the group.
- MLD report or done messages received on a host interface are forwarded to multicast-router interfaces in the same VLAN, but not to other host interfaces in the VLAN.
- Proxy MLD membership reports received with a null source IP address are accepted, to support report suppression.
- All unrecognized MLD packets are flooded to all (STP) unblocked member ports of the VLAN, to ensure that no data traffic is black-holed.

Data forwarding rules ensure that the multicast traffic received at the switch is forwarded to all interested downstream port members. Forwarding rules can be based on either Layer 3 multicast destination IP group address or Layer 2 destination MAC address.

- If a switch is already in a learned multicast group, multicast packets are forwarded only to those host interfaces in the VLAN that are members of the multicast group and to all multicast-router interfaces in the VLAN.
- If a switch is not in a learned multicast group, multicast packets for a group that has no current members are flooded to all member ports of the VLAN, as well as to all multicast-router interfaces in the VLAN. This behavior depends on whether the restrict-unknown-multicast feature, available only for multicast profiles, is enabled or not. (The default behavior is to flood packets on all ports. Refer to Restricting unknown multicast on page 39.) When it is enabled, multicast packets for a group that has no current members are forwarded to all multicast-router interfaces in the VLAN.

  NOTE
  For this release, Brocade VDX devices use Layer 2 multicast destination-MAC-address-based forwarding.

MLD snooping supports the following scale numbers:

TABLE 12 MLD snooping scale numbers

Parameter	Maximum
Number of IPv6 multicast snooping flows	4000
Number of VLANs	256
Groups learning rate	512/sec

The remainder of this section presents the tasks related to MLD configuration that are supported in this release.

Enabling and disabling MLD snooping globally

NOTE
The global and interface (VLAN) configurations of IPv6 MLDv1 Layer 2 snooping are independent of each other. However, MLD snooping must first be enabled globally for it to be enabled on a VLAN. (By default, snooping is disabled on VLANs.) If MLD snooping is disabled globally, the VLAN-level snooping configurations are retained in the running configuration but their functionality is disabled.

Do the following to enable and disable MLD snooping globally, respectively.

1. Enable MLD snooping globally.

```
device(config)# ipv6 mld snooping enable
```

2. Disable MLD snooping globally.

```
device(config)# no ipv6 mld snooping enable
```

Enabling and disabling MLD snooping at the interface level

NOTE

The global and interface (VLAN) configurations of IPv6 MLDv1 Layer 2 snooping are independent of each other. However, MLD snooping must first be enabled globally for it to be enabled on a VLAN. (By default, snooping is disabled on VLANs.) If MLD snooping is disabled globally, the VLAN-level snooping configurations are retained in the running configuration but their functionality is disabled.

Do the following to enable and disable MLD snooping on a VLAN, respectively.

1. Enable MLD snooping on a VLAN.

```
device(config)# int vlan 2000
device(config-Vlan-2000)# ipv6 mld snooping enable
```

2. Disable MLD snooping on a VLAN.

```
device(config)# int vlan 2000
device(config-Vlan-2000)# no ipv6 mld snooping enable
```

Enabling and disabling MLD querier functionality on a VLAN

You can use the MLD querier functionality to support MLD snooping on a VLAN where PIM and MLD are not enables (for example, because multicast traffic does not need to be routed). MLD querier functionality is disabled by default.

To enable this functionality, use the **ipv6 mld snooping querier enable** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# ipv6 mld snooping querier enable
```

To disable this functionality, use the **no ipv6 mld snooping querier enable** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# no ipv6 mld snooping querier enable
```

Configuring and unconfiguring an MLD static group on a VLAN

You can forward traffic statically for a multicast group onto a specified interface, so that the interface behaves as if MLD were enabled.

To enable this functionality, use the **ipv6 mld static-group** command on a VLAN interface, then select a multicast address to be joined, as well as a physical interface, as in the following example:

```
device(config-Vlan-2000)# ipv6 mld static-group ff1e::1 int te 54/0/1
```

To disable this functionality, use the **no ipv6 mld static-group** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# no ipv6 mld static-group ff1e::1 int te 54/0/1
```

Enabling and disabling MLD fast-leave on a VLAN

MLD fast-leave allows a group entry to be removed immediately from the receiver as soon as a done message is received, as long as the receiver is the only one on the segment that is subscribed to a group. This minimizes the leave latency of group memberships on an interface, as the device does not send group-specific queries. As a result, the group entry is removed from the multicast forwarding table as soon as a group done (leave) message is received.

> NOTE
> Use this command only if there is one receiver behind the interface for a given group.

Use the **ipv6 mld snooping fast-leave** command on a VLAN interface to enable this feature, as in the following example.

```
device(config-Vlan-2000)# ipv6 mld snooping fast-leave
```

Use the **no ipv6 mld snooping fast-leave** command on a VLAN interface to disable this feature, as in the following example.

```
device(config-Vlan-2000)# no ipv6 mld snooping fast-leave
```

Configuring the MLD query interval

You can configure the frequency at which MLD host query messages are sent. Larger values cause queries to be sent less often.

To set the MLD query interval, use the **ipv6 mld query-interval** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# ipv6 mld query-interval 1200
```

> NOTE
> The value set by this command must be greater than the query maximum response time, set by the **ipv6 mld query-max-response-time** command. Refer to the *Brocade Network OS Command Reference* for all ranges and defaults for the commands in this section.

To restore the default value, use the **no ipv6 mld query-interval** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# no ipv6 mld query-interval
```

Configuring the MLD last-member query interval

You can set the frequency at which MLD last-member query messages are sent. This is the interval for the response to a query sent after a host leave message is received from the last known active host on the subnet. The group is deleted if no reports are received in this interval. This interval adjusts the speed at which messages are transmited on the subnet. Smaller values detect the loss of a group member faster.

NOTE
If this interval is not configured explicitly, the value is taken from the robustness variable.

To set the MLD last-member query interval, use the **ipv6 mld last-member-query-interval** command on a VLAN interface, as in the following example:

```
switch(config-Vlan-2000)# ipv6 mld last-member-query-interval 1500
```

To restore the default value, use the **no ipv6 mld last-member-query-interval** command on a VLAN interface, as in the following example:

```
switch(config-Vlan-2000)# no ipv6 mld last-member-query-interval
```

Configuring the MLD last-member query count

You can set the number of times that an MLD query is sent in response to a host leave message. This is the number of times, separated by the last-member query-response interval (configured by the **ipv6 mld last-member-query-interval** command), that an MLD query is sent in response to a host leave message from the last known active host on the subnet.

NOTE
If this interval is not configured explicitly, the value is taken from the robustness variable.

To change the MLD last-member query count from the default, use the **ipv6 mld last-member query count** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# ipv6 mld last-member-query-count 3
```

To restore the default value, use the **no ipv6 mld last-member-query-count** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# no ipv6 mld last-member-query-count
```

Configuring the MLD query maximum response time

You can configure the maximum response time for IPv6 MLDv1 snooping MLD queries for a specific VLAN interface, as in the following example:

```
device(config)# int vlan 2000
device(config Vlan-2000)# ipv6 mld query-max-response-time 15
```

NOTE
Larger values spread out host responses over a longer time. The value set by this command must be less than the general query interval, set by the **ipv6 mld query-interval** command.

To restore the default value, use the **no ipv6 mld query-max-response-time** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# no ipv6 mld query-max-response-time
```

Configuring the MLD snooping robustness variable

A robustness value can be configured to compensate for packet loss in congested networks. This value determines the number of general MLD snooping queries that are sent before a multicast address is aged out for lack of a response. The default is 2.

To change the default robustness variable on a VLAN, use the **ipv6 mld snooping robustness-variable** command, as in the following example:

```
switch(config-Vlan-2000)# ipv6 mld snooping robustness-variable 7
```

To restore the default value, use the **no ipv6 mld robustness-variable** command on a VLAN interface, as in the following example:

```
switch(config-Vlan-2000)# no ipv6 mld robustness-variable
```

Configuring the MLD startup query count

The IPv6 MLDv1 startup query count is the number of queries that are separated by the startup interval. The default is 1.

Do the following to change the startup-query interval on a VLAN, as in the following example.

```
device(config-Vlan-2000)# ipv6 mld startup-query-count 2
```

To restore the default value, use the **no ipv6 mld startup-query-count** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# no ipv6 mld startup-query-count
```

Configuring the MLD startup query interval

You can change the query interval between the general queries that are sent by the querier on startup. The default interval is 1. The querier may be the MLD snooping querier or an external querier.

Do the following to change the startup-query interval on a VLAN, as in the following example.

```
device(config-Vlan-2000)# ipv6 mld startup-query-interval 2
```

To restore the default value, use the **no ipv6 mld startup-query-interval** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# no ipv6 mld startup-query-interval
```

Configuring a VLAN port member to be a multicast router port

You can configure a VLAN port member to be a multicast router (mrouter) port.

To configure a VLAN port member to be a multicast router (mrouter) port., use the **ipv6 mld snooping mrouter interface** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# ipv6 mld snooping mrouter interface te 54/0/1
```

To disable the VLAN port member from being an mrouter port., use the **ipv6 mld snooping mrouter interface** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# no ipv6 mld snooping mrouter interface te 54/0/1
```

Managing the flooding of multicast data traffic

You can deactivate or reactivate on a VLAN the flooding of unregistered multicast data traffic on IPv6 MLDv1 snooping-enabled VLANs.

To deactivate the flooding of unregistered multicast data traffic, use the **ipv6 mld snooping restrict-unknown-multicast** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# ipv6 mld snooping restrict-unknown-multicast
```

To reactivate the flooding of unregistered multicast data traffic, use the **no ipv6 snooping restrict-unknown-multicast** command on a VLAN interface, as in the following example:

```
device(config-Vlan-2000)# no ipv6 mld snooping restrict-unknown-multicast
```

Monitoring and managing MLD snooping

You can monitor MLD snooping by using a variety of **show** commands. In addition, you can clear the data for MLD groups and statistics by using **clear** commands. A **debug** command is also available. For command details, refer to the *Network OS Command Reference*.

The following table lists the available **show** commands for MLD snooping.

TABLE 13 MLD snooping show commands

Command	Description
show ipv6 mld groups	Displays information about IPv6 MLDv1 groups.
show ipv6 mld interface vlan	Displays IPv6 MLD information for a VLAN.
show ipv6 mld snooping	Displays IPv6 MLD snooping details.
show ipv6 mld statistics	Displays IPv6 MLDv1 statistics.

The following table lists the available **clear** and **debug** commands.

TABLE 14 MLD snooping clear and debug commands

Command	Description
clear ipv6 mld groups	Clears IPv6 MLDv1 group cache entries.
clear ipv6 mld statistics	Clears IPv6 MLDv1 snooping statistics.
debug ipv6 mld	Displays information related to IPv6 MLD, with a variety of options.