

PIM-DM (Dense Mode) on the 5300xl Switches

Contents

Overview	5-2
Introduction	5-3
Feature Overview	5-4
PIM-DM Operation	5-4
Multicast Flow Management	5-7
General Configuration Elements	5-9
Terminology	5-9
PIM-DM Operating Rules	5-10
Configuring PIM-DM on the Series 5300xl Switches	5-11
PIM Global Configuration Context	5-12
PIM VLAN (Interface) Configuration Context	5-15
Displaying PIM Data and Configuration Settings on the Series 5300xl Switches	5-22
Displaying PIM Route Data	5-23
Displaying PIM Status	5-27
Operating Notes	5-34
Troubleshooting	5-36
Messages Related to PIM Operation	5-37
Applicable RFCs	5-40
Exceptions to Support for RFC 2932 - Multicast Routing MIB ..	5-41

Overview

This chapter describes protocol-independent multicast routing operation on the ProCurve Series 5300xl switches and how to configure it with the switch's built-in interfaces, and assumes an understanding of multimedia traffic control with IP multicast (IGMP), which is described in chapter 4, "Multimedia Traffic Control with IP Multicast (IGMP)".

For general information on how to use the switch's built-in interfaces, refer to these chapters in the *Management and Configuration Guide* for your switch:

- Chapter 3, "Using the Menu Interface"
- Chapter 4, "Using the Command Line Interface (CLI)"
- Chapter 5, "Using the Web Browser Interface"
- Chapter 6, "Switch Memory and Configuration"

Introduction

This feature operates only on the Series 5300xl switches.

Feature	Default	Menu	CLI	Web
Configure PIM Global	n/a	—	5-12	—
Configure PIM VLAN Interface	n/a	—	5-15	—
Display PIM Route Data	Disabled	—	5-23	—
Display PIM Status	0 (Forward All)	—	5-27	—

In a network where IP multicast traffic is transmitted for multimedia applications, multicast traffic is blocked at routed interface (VLAN) boundaries unless a multicast routing protocol is running. PIM-DM (Protocol Independent Multicasting-Dense Mode, draft version 3) enables and controls multicast traffic routing on the Switch Series 5300XL devices.

PIM-DM is used in networks where, at any given time, multicast group members exist in relatively large numbers and are present in most subnets. PIM-DM operates with any unicast IPv4 routing protocol available on the switch. However, note that PIM-DM uses “flooding” to initially propagate a multicast group to a network, then prunes back the branches that have no hosts requiring membership. For this reason, ample bandwidth is a requirement in PIM-DM applications.

IGMP provides the communication link between a host and the multicast router running PIM. Where hosts are connected directly to the routing device, such as the 5300XL, both PIM and IGMP must run on the switch.

Feature Overview

PIM-DM on the Switch Series 5300XL devices includes:

- **Routing Protocol Support:** PIM uses whichever unicast routing protocol is running on the routing switch. These can include:
 - RIP
 - OSPF
 - Static routes
 - Directly connected interfaces
- **Interface Support:** PIM-DM supports up to 127 outbound VLANs (and 1 inbound VLAN) in its multicast routing table (MRT) at any given time, meaning that the sum of all outbound VLANs across all current flows on a routing switch may not exceed 63. (A single flow may span one inbound VLAN and up to 63 outbound VLANs, depending on the VLAN memberships of the hosts actively belonging to the flow.)
- **IGMP Compatibility:** PIM-DM is compatible with IGMP versions 1 - 3, and is fully interoperable with IGMP for determining multicast flows.
- **XRRP:** PIM-DM is fully interoperable with XRRP to quickly transition multicast routes in the event of a failover.
- **MIB Support:** With some exceptions, PIM-DM supports the parts of the Multicast Routing MIB applicable to PIM-DM operation. (Refer to “Exceptions to Support for RFC 2932 - Multicast Routing MIB” on page 5-41.)
- **PIM Draft Specifications:** Compatible with PIM-DM draft specification, versions 1 and 2.

PIM-DM Operation

PIM-DM operates at the router level to direct traffic for a particular multicast group along the most efficient path to the VLANs having hosts that have joined that group. A unicast source address and a multicast group address comprise a given source/group (S/G) pair. Multicast traffic moving from a source to a multicast group address creates a *flow* to the area(s) of the network requiring the traffic. That is, the flow destination is the multicast group address, and not a specific host or VLAN. Thus, a single multicast flow has one source and one

multicast group address (destination), but may reach many hosts in different subnets, depending on which hosts have issued joins for the same multicast group.

PIM routes the multicast traffic for a particular S/G pair on paths between the source unicast address and the VLANs where it is requested (by joins from hosts connected to those VLANs). Physical destinations for a particular multicast group can be hosts in different VLANs or networks. Individual hosts use IGMP configured per-VLAN to send joins requesting membership in a particular multicast group. All hosts that have joined a given multicast group (defined by a multicast address) remain in that group as long as they continue to issue periodic joins.

On the Switch Series 5300XL devices, PIM-DM interoperates with IGMP and the switch's routing protocols. (Note that PIM-DM operates independently of the routing protocol you choose to run on your switches, meaning you can use PIM-DM with RIP, OSPF, or static routes configured.) PIM-DM utilizes a unicast routing table to find the path to the originator of the multicast traffic and sets up multicast "trees" for distributing multicast traffic. (This method is termed *reverse path forwarding*, or *RPF*).

For the flow of a given multicast group, PIM-DM creates a "tree" structure between the source and the VLANs where hosts have joined the group. The tree structure consists of:

- Extended branches to VLANs with hosts that currently belong to the group
- Pruned branches to VLANs with no hosts that belong to the group

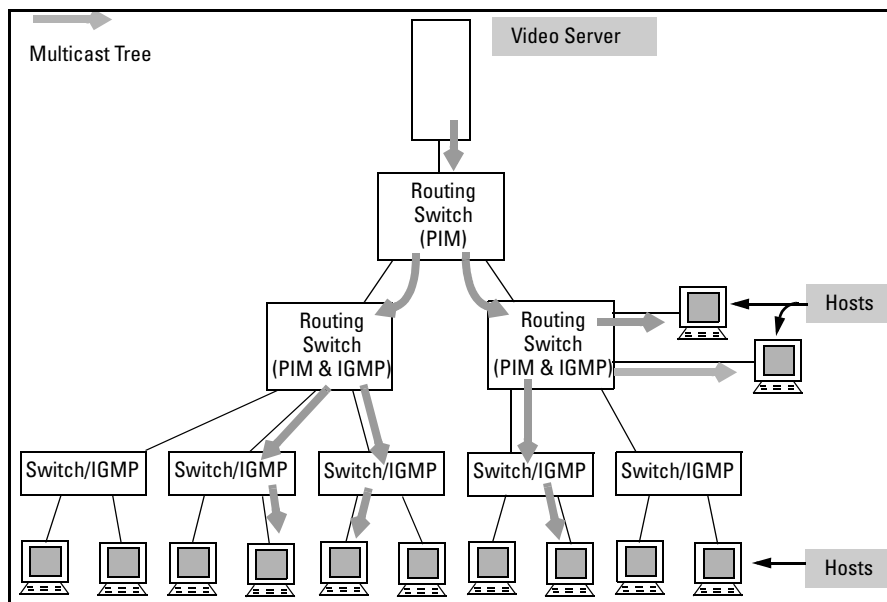


Figure 5-1. Example of Multicast “Tree” for a Given Flow

When the routing switch detects a new multicast flow, it initially floods the traffic throughout the PIM-DM domain, and then prunes the traffic on the branches (network paths) where joins have not been received from individual hosts. This creates the “tree” structure shown above. The routing switch maintains individual branches in the multicast tree as long as there is at least one host maintaining a membership in the multicast group. When all of the hosts in a particular VLAN drop out of the group, PIM-DM prunes that VLAN from the multicast tree. Similarly, if the routing switch detects a join from a host in a pruned VLAN, it adds that branch back into the tree.

Note

Where the multicast routers in a network use one or more multinetted VLANs, there must be at least one subnet common to all routers on the VLAN. This is necessary to provide a continuous forwarding path for the multicast traffic on the VLAN. Refer to the [**all | source-ip-address**] option under “PIM VLAN (Interface) Configuration Context” on page 5-15.

Multicast Flow Management

This section provides details on how the routing switch manages forwarding and pruned flows. This information is useful when planning topologies to include multicast support and when viewing and interpreting the “show” command output for PIM-DM features.

Initial Flood and Prune. As mentioned earlier, when a router running PIM-DM receives a new multicast flow, it initially floods the traffic to all downstream multicast routers. PIM-DM then prunes the traffic on paths to VLANs that have no host joins for that multicast address. (Note that PIM-DM does not re-forward traffic back to its source VLAN.)

Maintaining the Prune State. For a multicast group “X” on a given VLAN, when the last host belonging to group “X” leaves the group, PIM places that VLAN in a prune state, meaning the group “X” multicast traffic is blocked to that VLAN. The prune state remains until a host on the same VLAN issues a join for group “X”, in which case the router cancels the prune state and changes the flow to the forwarding state.

State Refresh Packets and Bandwidth Conservation. A 5300XL multicast router, if directly connected to a multicast source such as a video conferencing application, periodically transmits *state refresh* packets to downstream multicast routers. On routers that have pruned the multicast flow, the state refresh packets keep the pruned state alive. On routers that have been added to the network after the initial flooding and pruning of a multicast group, the state refresh packets inform the newly added router of the current state of that branch. This means that if all multicast routers in a network support the state refresh packet, then the multicast router directly connected to the multicast source performs only one flood-prune cycle to the edge of the network when a new flow (multicast group) is introduced, and preserves bandwidth for other uses. Note, however, that some vendors’ multicast routers do not offer the state refresh feature. In this case, PIM-DM must periodically advertise an active multicast group to these devices by repeating the flood/prune cycle on the paths to such routers.. For better traffic management in multicast-intensive networks where some multicast routers do not offer the state refresh feature, you may want to group such routers where the increased bandwidth usage will have the least effect on overall network performance.

PIM-DM (Dense Mode) on the 5300xl Switches
PIM-DM Operation

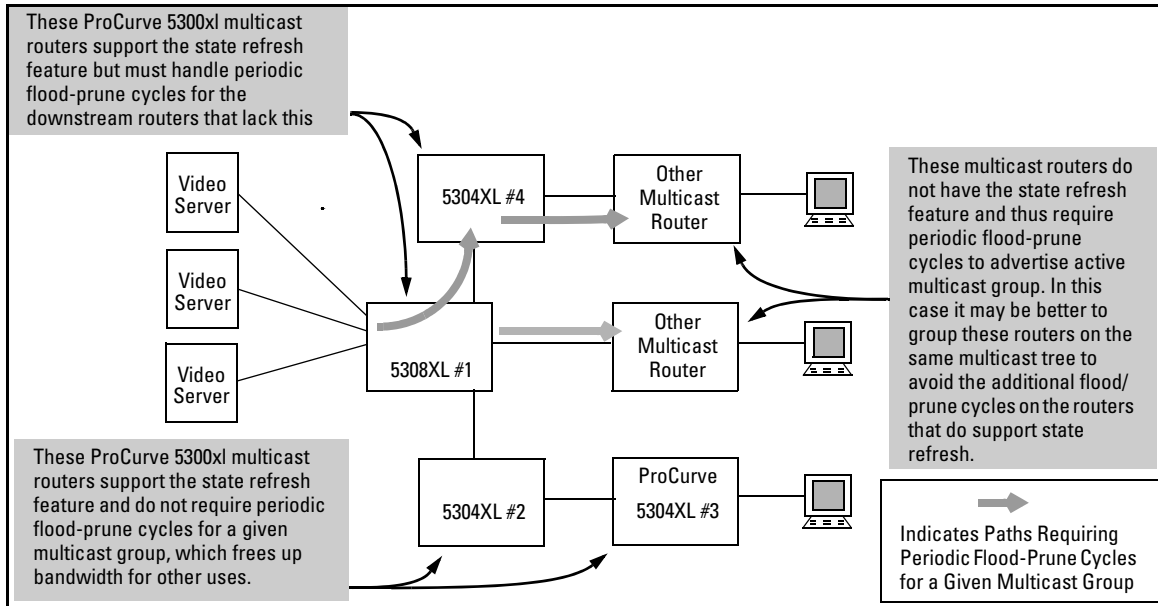


Figure 5-2. Example of Bandwidth Conservation in 5300XL Routing Switches with PIM-DM State Refresh

General Configuration Elements

The configured elements PIM-DM requires are:

1. IP routing enabled on all routing switches you want to carry routed multicast traffic.
2. Configure the routing method(s) needed to reach the interfaces (VLANs) on which you want multicast traffic available for hosts in your network:
 - Enable RIP or OSPF at both the global and VLAN levels on the routers where there are connected hosts that may issue multicast joins.
 - Configure static routes to and from the destination subnets.
3. Enable IP multicast routing.
4. For each VLAN on which there are hosts that you want to join multicast groups, enable IGMP on that VLAN. Repeat this action on every switch and router belonging to the VLAN.
5. Enable PIM-DM at the global level on the routing switch and on the VLANs where you want to allow routed multicast traffic.

Note

When you initially enable PIM-DM, HP recommends that you leave the PIM-DM configuration parameters at their default settings. You can then assess performance and make configuration changes where a need appears.

Terminology

Flow: Multicast traffic moving between a unicast source and a multicast group. One S/G pair is counted as a single flow, regardless of the number of hosts belonging to the related multicast group.

Host: A client device that requests multicast traffic by transmitting IGMP “joins” for a specific multicast group, such as a video conferencing application.

MRT (Multicast Routing Table). The routing switch creates this table internally to maintain data on each multicast group it supports. The “Show” commands described later in this chapter display MRT data managed in this table.

Multicast Address: In IP multicast traffic on the switch, this is a single IP address that can be used by a group of related or unrelated clients wanting the same data. A single S/G pair consists of unicast source address and a multicast group address. Sometimes termed a “multicast group address”. See also “Source” and “S/G Pair”.

Multicast Routing: A method for transmitting multicast datagrams from a source in one IP network to a multicast address in one or more other IP networks.

PIM Neighbor: On a routing switch configured for PIM operation, a PIM neighbor is another PIM-configured routing switch or router that is either directly connected to the first routing switch or connected through networked switches and/or hubs.

Prune: To eliminate branches of a multicast tree that have no hosts sending joins to request or maintain membership in that particular multicast group.

S/G Pair: The unicast address of the server transmitting the multicast traffic and the multicast address to which the server is transmitting the traffic.

Source (S): In IP multicast traffic on the switch, the source (S) is the unicast address of the server transmitting the multicast traffic. A single S/G pair consists of unicast source address and a multicast group address. See also “S/G Pair”.

PIM-DM Operating Rules

- The routing switch supports 1022 multicast flows in hardware and up to 978 additional flows in software. (For more on this topic, refer to “Flow Capacity” on page 5-35.)
- The multicast routing table (MRT) that PIM-DM creates allows up to 127 outbound VLANs, meaning that at any given time, PIM-DM supports multicast routing across 127 VLANs.
- The routing switch allows one instance of PIM per VLAN. Thus, in networks using multinetted VLANs, all routers on a given VLAN intended to route multicast packets must have a least one common subnet on that VLAN. Thus, in the case of multinetting, you must select one subnet on the multinetted VLAN to use for multicast routing. To facilitate this, the routing switch provides a command for specifying which IP address PIM will use on each VLAN.

Configuring PIM-DM on the Series 5300xl Switches

Command	Page
PIM Global Context Commands	
[no] ip multicast-routing	5-12
[no] router pim	5-12
state-refresh	5-12
trap	5-13
PIM Interface Context Commands	
[no] ip pim	5-15
[< all source-ip-address >]	5-15
[hello-interval]	5-15
[hello-delay]	5-16
[graft-retry-interval]	5-16
[max-graft-retries]	5-17
[lan-prune-delay]	5-17
[propagation-delay]	5-18
[override-delay]	5-18
[ttl-threshold]	5-19

PIM-DM requires configuration on both the global level and on the VLAN (interface) level. The recommended configuration order is:

1. Enable IGMP on all VLANs where hosts may join a multicast group.
2. Enable the following at the global level on the Switch Series 5300XL device.
 - IP routing
 - IP multicast routing
 - Router PIM and any non-default, global PIM settings you want to apply
 - Router RIP, Router OSPF, and/or a static route
3. If you selected RIP or OSPF in step step 2, then on each VLAN where you want multicast routing to operate, enable the same option.
4. Enable the following in each VLAN context where you want multicast routing to operate:
 - IP RIP or IP OSPF
 - IP PIM
 - Any non-default, VLAN-level IP PIM settings you want to apply

PIM Global Configuration Context

Note

PIM-DM operation requires a routing protocol enabled on the routing switch. You can use RIP, OSPF, and/or static routing. The examples in this section use RIP. For more on these topics, refer to chapter 11, “IP Routing Features”, in this guide.

Syntax: [no] ip multicast-routing

Enables or disables IP multicast routing on the routing switch. IP routing must be enabled. (Default: Disabled.)

Syntax: [no] router pim

Enables or disables PIM at the global level. IP routing must be enabled first. (Default: Disabled.)

Syntax: router pim [state-refresh < 10 - 300 >]

Sets the interval in seconds between successive State Refresh messages originated by the routing switch. Note that only the routing switch connected directly to the unicast source initiates state-refresh packets. All other PIM routers in the network only propagate these state-refresh packets. (Default: 60 seconds)

Syntax: [no] router pim trap < all | neighbor-loss | hardware-mrt-full | software-mrt-full >

Enables and disables these PIM SNMP traps:

all — *Enable/Disable all PIM notification traps.*

neighbor-loss — *Enable/Disable the notification trap sent when the timer for a multicast router neighbor expires and the switch has no other multicast router neighbors on the same VLAN with a lower IP address. (Default: Disabled.)*

hardware-mrt-full — *Enable/Disable notification trap when the hardware multicast routing table (MRT) is full (1023 active flows). In this state, any additional flows are handled by the software MRT, which increases processing time for the affected flows. (Default: Disabled.)*

software-mrt-full — *Enable/Disable notification trap when the routing switch's software multicast routing table is full (that is, when routing resources for active flows are exhausted). (Default: Disabled.) Note that in this state, the routing switch does not accept any additional flows.*

Example of Configuring PIM at the Global Level. In figure 5-1 on page 5-6, the “5308XL #1” routing switch is directly connected to the multicast sources for the network. In this case, suppose that you want to do the following:

- Reduce the state-refresh time from the default 60 seconds to 30 seconds. Note that the routing switch transmits state-refresh packets only if it is directly connected to the multicast source.
- Configure an SNMP trap to notify your network management station if the routing switch's hardware multicast routing table becomes filled to the maximum of 1023 active flows.

To configure global-level PIM operation for the “5308XL #1” routing switch, you would use the commands shown in figure 5-3, below.

```
ProCurve(config)# ip routing
ProCurve(config)# ip multicast-routing
ProCurve(config)# router pim
ProCurve(pim)# router rip
ProCurve(pim)# state-refresh 45
ProCurve(pim)# trap hardware-mrt-full
ProCurve(pim)# write mem
ProCurve(pim)# exit
```

← Enables IP routing.
← Enables multicast routing.
← Enables PIM.
← Enables RIP.
← Configures a non-default State Refresh timer.
← Sets an SNMP trap to notify an SNMP management station if the hardware multicast routing table fills with active flows.

Using show config displays the configuration changes resulting from the above commands.

```
ProCurve(config)# show config
Startup configuration:
; J4850A Configuration Editor; Created on release #E.08.01
hostname "HPswitch"
module 1 type J4820A
ip routing
snmp-server community "public" Unrestricted
snmp-server host 15.29.38.205 "public" Not-INFO
snmp-server host 15.255.124.84 "public" Not-INFO
vlan 1
  name "DEFAULT_VLAN"
  untagged A1-A5,A16,A19
  ip address dhcp-bootp
  no untagged A6-A15,A17-A18,A20-A24
  exit
vlan 29
  name "VLAN29"
  :
  :
  exit
vlan 25
  name "VLAN25"
  untagged A20-A24
  ip address 25.38.10.1
  exit
ip multicast-routing
ip ssh filetransfer
router rip
  exit
router pim
  trap hardware-mrt-full
  state-refresh 45
  exit
```

←
←
←
←

Figure 5-3. Example of Configuring PIM-DM on a Routing Switch at the Global Level

After configuring the global-level PIM operation on a routing switch, go to the device's VLAN context level for each VLAN you want to include in your multicast routing domain. (Refer to "PIM VLAN (Interface) Configuration Context", below.

PIM VLAN (Interface) Configuration Context

Syntax: [no] ip pim
[no] vlan < vid > ip pim

*Enables multicast routing on the VLAN interface to which the CLI is currently set. The **no** form disables PIM on the VLAN. Default: Disabled.*

Syntax: [no] ip pim [all | < source-ip-address >]
[no] vlan < vid > ip pim [all | < source-ip-address >]

*In networks using multinetted VLANs, all routers on a given VLAN intended to route multicast packets must have a least one common subnet on that VLAN. Use this command when the VLAN is configured with multiple IP addresses (multinetting) to specify the IP address to use as the source address for PIM protocol packets outbound on the VLAN. Use < ip-address > to designate a single subnet in cases where multicast routers on the same multinetted VLAN are not configured with identical sets of subnet IP addresses . Use < all > if the multinetted VLAN is configured with the same set of subnet addresses. (**Default:** The Primary VLAN.)*

Syntax: ip pim [hello-interval < 5 - 30 >]
vlan < vid > ip pim [hello-interval < 5 - 30 >]

*Changes the frequency at which the routing switch transmits PIM "Hello" messages on the current VLAN. The routing switch uses "Hello" packets to inform neighboring routers of its presence. The routing switch also uses this setting to compute the **Hello Hold Time**, which is included in Hello packets sent to neighbor routers. **Hello Hold Time** tells neighbor routers how long to wait for the next Hello packet from the routing switch. If another packet does not arrive within that time, the router removes the neighbor adjacency on that VLAN from the routing table, which removes any flows running on that interface. Shortening the Hello interval reduces the Hello Hold Time. This has the effect of changing how quickly other routers will stop sending traffic to the routing switch if they do not receive a new Hello packet when expected.*

*For example, if multiple routers are connected to the same VLAN and the routing switch requests multicast traffic, all routers on the VLAN receive that traffic. (Those which have pruned the traffic will drop it when they receive it.) If the upstream router loses contact with the routing switch receiving the multicast traffic (that is, fails to receive a Hello packet when expected), then the shorter Hello Interval causes it to stop transmitting multicast traffic onto the VLAN sooner, resulting in less unnecessary bandwidth usage. Not used with the **no** form of the **ip pim** command.*

Syntax: ip pim [hello-delay < 0 - 5 >]
vlan < vid > ip pim [hello-delay < 0 - 5 >]

*Changes the maximum time in seconds before the routing switch actually transmits the initial PIM Hello message on the current VLAN. In cases where a new VLAN activates with connections to multiple routers, if all of the connected routers sent Hello packets at the same time, then the receiving router could become momentarily overloaded. This value randomizes the transmission delay to a time between **0** and the **hello delay** setting. Using “**0**” means no delay. After the routing switch sends the initial Hello Packet to a newly detected VLAN interface, it sends subsequent Hello packets according to the current **Hello Interval** setting. Not used with the **no** form of the **ip pim** command. Default: 5 seconds.*

Syntax: ip pim [graft-retry-interval < 1-10 >]
vlan < vid > ip pim [graft-retry-interval < 1-10 >]

*Graft packets result when a downstream router transmits a request to join a flow. The upstream router responds with a graft acknowledgment packet. If the Graft Ack is not received within the time period of the **graft-retry-interval**, it resends the graft packet. This command changes the interval (in seconds) the routing switch waits for the Graft Ack (acknowledgement) from another router before resending the Graft request. Not used with the **no** form of the **ip pim** command. (Default: 3 seconds.)*

Syntax: ip pim [max-graft-retries < 1 - 10 >
vlan < vid > ip pim [max-graft-retries < 1 - 10 >

*Changes the number of times the routing switch will retry sending the same graft packet to join a flow. If a Graft Ack response is not received after the specified number of retries, the routing switch ceases trying to join the flow. In this case the flow is removed until either a state refresh from upstream re-initiates the flow or an upstream router floods the flow. Increasing this value helps to improve multicast reliability. Not used with the **no** form of the **ip pim** command. (Default: 3 attempts.)*

Syntax: ip pim [lan-prune-delay]
vlan < vid > ip pim [lan-prune-delay]

*Enables the LAN Prune Delay option on the current VLAN. With **lan-prune-delay** enabled, the routing switch informs downstream neighbors how long it will wait before pruning a flow after receiving a prune request. Other, downstream routers on the same VLAN must send a Join to override the prune before the **lan-prune-delay** time if they want the flow to continue. This prompts any downstream neighbors with hosts continuing to belong to the flow to reply with a Join. If no joins are received after the **lan-prune-delay** period, the routing switch prunes the flow. The **propagation-delay** and **override-interval** settings (below) determine the **lan-prune-delay** setting.*

*Uses the **no** form of the **ip pim** command to disable the LAN Prune Delay option. (Default: Enabled.)*

Syntax: ip pim [propagation-delay < 250-2000 >]
vlan < vid > ip pim [propagation-delay < 250-2000 >]

ip pim [override-interval < 500 - 6000 >]
vlan < vid > ip pim [override-interval < 500 - 6000 >]

*A routing switch sharing a VLAN with other multicast routers uses these two values to compute the **lan-prune-delay** setting (above) for how long to wait for a PIM-DM join after receiving a prune packet from downstream for a particular multicast group. For example, a network may have multiple routing switches sharing VLAN “X”. When an upstream routing switch initially floods traffic from multicast group “X” to VLAN “Y”, if one of the routing switches on VLAN “Y” does not want this traffic it issues a prune response to the upstream neighbor. The upstream neighbor then goes into a “prune pending” state for group “X” on VLAN “Y”. (During this period, the upstream neighbor continues to forward the traffic.) During the “pending” period, another routing switch on VLAN “Y” can send a group “X” Join to the upstream neighbor. If this happens, the upstream neighbor drops the “prune pending” state and continues forwarding the traffic. But if no routers on the VLAN send a Join, then the upstream router prunes group “X” from VLAN “Y” when the **lan-prune-delay** timer expires. (Defaults: **propagation-delay** = 500 milliseconds; **override-interval** = 2500 milliseconds.)*

Syntax: ip pim [ttl-threshold < 0 - 255 >]
vlan < vid > ip pim

Sets the multicast datagram time-to-live (router hop-count) threshold for the VLAN. Any IP multicast datagrams or state refresh packets with a TTL less than this threshold will not be forwarded out the interface. The default value of 0 means all multicast packets are forwarded out the interface.

*This parameter provides a method for containing multicast traffic within a network, or even within specific areas of a network. Initially, the multicast traffic source sets a TTL value in the packets it transmits. Each time one of these packets passes through a multicast routing device, the TTL setting decrements by 1. On a Switch Series 5300XL device, if the packet arrives with a TTL lower than the **mroute ttl-threshold**, the routing switch does not forward the packet. Changing this parameter on a routing switch requires knowledge of the TTL setting of incoming multicast packets. A value that is too high can allow multicast traffic to go beyond your internal network. A value that is too low may prevent some intended hosts from receiving the desired multicast traffic. (Default: 0 — forwards multicast traffic regardless of packet TTL setting.)*

Example of Configuring PIM-DM Operation at the VLAN Level. The network in figure 5-4 uses VLAN 25 for multicast traffic. However, this VLAN is multinetted and there is only one subnet (25.38.10.x) in VLAN 25 that is common to all three routing switches. Thus, when configuring VLAN 25 on these routing switches to perform multicast routing, it is necessary to use **ip pim < source-ip-address >** to designate the common subnet as the source address for outbound multicast traffic on VLAN 25. (If only identical subnets were present in the multinetted VLAN 25 configuration on all three devices, then the **ip pim all** command would be used instead.) Note that the other VLANs in the network are not multinetted and therefore do not require the **ip pim < all | source-ip-address >** option.

For this example, assume that the VLANs and IP addressing are already configured on the routing switch.

PIM-DM (Dense Mode) on the 5300xl Switches
 Configuring PIM-DM on the Series 5300xl Switches

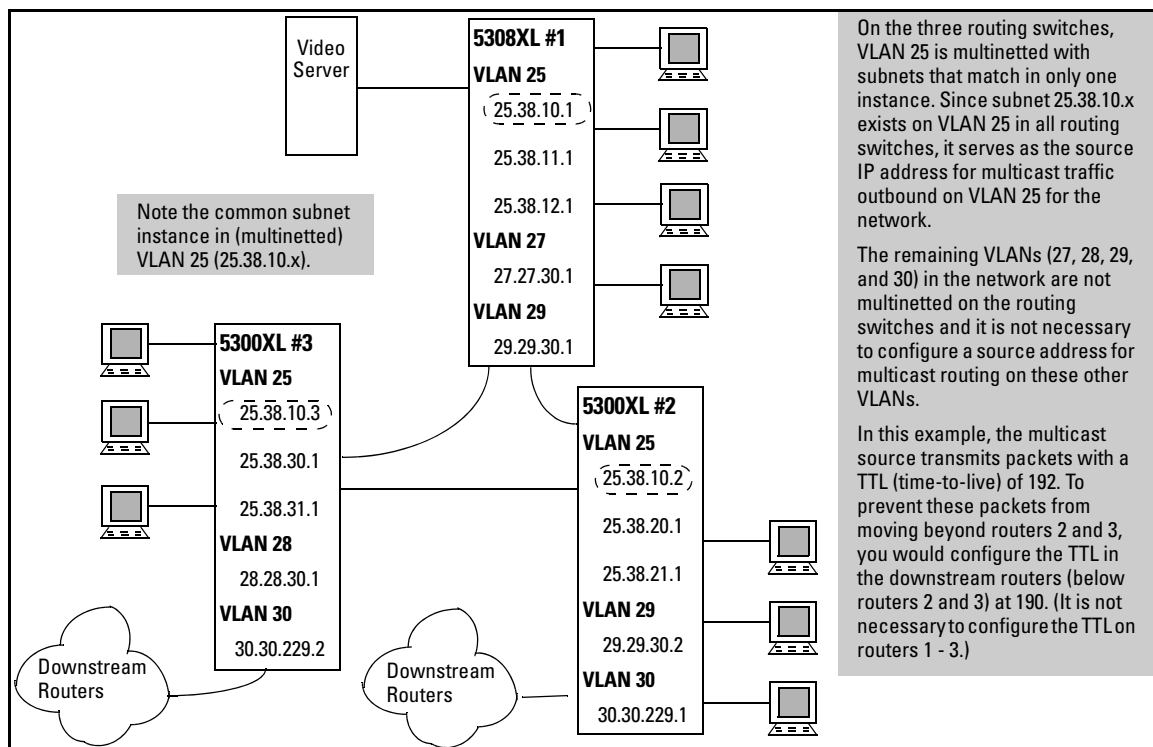


Figure 5-4. Example of a Multicast Network with a Multinetted VLAN

Figure 5-5 illustrates the configuration steps for configuring multicast routing at the VLAN level for the 5300XL #1 routing switch shown in figure 5-4.

```

ProCurve(config)# vlan 25
ProCurve(vlan-25)# ip igmp
ProCurve(vlan-25)# ip rip
ProCurve(vlan-25)# ip pim 25.38.10.1
ProCurve(vlan-25)# vlan 27

ProCurve(vlan-27)# ip igmp
ProCurve(vlan-27)# ip rip
ProCurve(vlan-27)# ip pim

ProCurve(vlan-27)# vlan 29
ProCurve(vlan-29)# ip igmp
ProCurve(vlan-29)# ip rip
ProCurve(vlan-29)# ip pim

ProCurve(vlan-29)# write mem
ProCurve(vlan-29)# exit
    
```

Figure 5-5. VLAN-Level Configuration Steps To Support PIM-DM on the 5308XL #1 Device

```

ProCurve(config)# show config
Startup configuration:
; J4850A Configuration Editor; Created on release #E.08.XX

hostname 'ProCurve'
module 1 type J4820A
ip_routing,
.
.
vlan 29
 name "VLAN29"
 untagged A11-A15,A17
 ip address 29.29.30.1 255.255.248.0
 ip igmp,
 exit
vlan 27
 name "VLAN27"
 untagged A6-A10,A18
 ip address 27.27.30.1 255.255.248.0
 ip igmp,
 exit
vlan 25
 name "VLAN25"
 untagged A20-A24
 ip address 25.38.10.1 255.255.255.0
 ip address 25.38.11.1 255.255.255.0
 ip address 25.38.12.1 255.255.255.0
 ip igmp,
 exit
ip multicast-routing
router rip
 exit
router pim
 trap hardware-mrt-full
 state-refresh 45
 exit
vlan 25
 ip rip
 ip pim 25.38.10.1
 ip pim ttl-threshold 2
 exit
vlan 27
 ip rip
 ip pim all
 ip pim ttl-threshold 2
 exit
vlan 29
 ip rip
 ip pim all
 ip pim ttl-threshold 2
 exit
    
```

Enables IP routing; required for multicast routing.

Multinetting and IGMP enabled in VLAN 25.

Multicast Routing Configuration for Global Level..

Indicates the source-IP-address for multicast packets forwarded on this VLAN.

Multicast Routing Configuration for VLAN 25.

Multicast Routing Configurations for VLANs 27 and 29.

Dashed lines indicate configuration settings affecting multicast routing.

Figure 5-6. The Configuration Supporting Multicast Routing on the 5300XL #1 Routing Switch Shown in Figure 5-4 (Page 5-20)

Displaying PIM Data and Configuration Settings on the Series 5300xl Switches

Command	Page
show ip mroute	5-23
[interface < vid >]	5-24
[< multicast-ip-addr > < source-ip-addr >]	5-25
show ip pim	5-27
[interface	5-28
[< vid >]]	5-29
[mroute	5-30
[< multicast-group-address > < multicast-source-address >]]	5-31
neighbor	5-33
[< ip-address >]	5-34

Displaying PIM Route Data

Syntax: show ip mroute

Without parameters, lists all VLANs actively forwarding routed, multicast traffic.

Group Address: *The multicast address of the specific multicast group (flow).*

Source Address: *The unicast address of the multicast group source.*

Neighbor: *The IP address of the upstream multicast router interface (VLAN) from which the multicast traffic is coming. A blank field for a given multicast group indicates that the multicast server is directly connected to the routing switch.*

VLAN: *The interface on which the multicast traffic is moving.*

For example, the next figure displays the show ip route output on the “5300XL #2” routing switch in figure 5-4 on page 5-20. This case illustrates two multicast groups from the same multicast server source.

```
ProCurve(config)# show ip mroute
```

Group Address	Source Address	Neighbor	VLAN
239.255.255.1	27.27.30.2	29.29.30.1	29
239.255.255.5	27.27.30.2	29.29.30.1	29

Indicates the upstream multicast router interface (VLAN) from which the multicast traffic is coming.

Figure 5-7. Example Showing the Route Entry Data on the “5300XL #2” Routing Switch in Figure 5-4 on Page 5-20

Syntax: show ip mroute [interface < vid >]

Lists these settings:

VLAN: *The VID specified in the command.*

Protocol Identity: *PIM-DM only.*

TTL: *The time-to-live threshold for packets forwarded through this VLAN. When configured, the routing switch drops multicast packets having a TTL lower than this value. (When a packet arrives, the routing switch decrements it's TTL by 1, then compares the decremented packet TTL to the value set by this command.) A TTL Threshold setting of 0 (the default) means all multicast packets are forwarded regardless of the TTL value they carry. A multicast packet must have a TTL greater than 1 when it arrives at the routing switch. Otherwise the routing switch drops the packet instead of forwarding it on the VLAN.*

```
ProCurve(config)# show ip mroute interface 29
IP Multicast Interface
VLAN      : 29
Protocol  : PIM-DM
TTL Threshold : 0
```

Figure 5-8. Example of the Above Command on the “5300XL #2” Routing Switch in Figure 5-4 on Page 5-20

Syntax: show ip mroute [< multicast-ip-addr > < source-ip-addr >]

Lists the following data for the specified flow (multicast group):

Group Address: *The multicast group IP address for the current group.*

Source Address: *The multicast source address < source-ip-addr > for the current group.*

Source Mask: *The subnet mask applied to the multicast source address < source-ip-addr >.*

Neighbor: *Lists the IP address of the upstream next-hop router running PIM-DM; that is, the router from which the routing switch is receiving datagrams for the current multicast group. This value is **0.0.0.0** if the routing switch has not detected the upstream next-hop router's IP address. This field is empty if the multicast server is directly connected to the routing switch.*

VLAN: *Lists the VLAN ID (VID) on which the routing switch received the specified multicast flow.*

Up Time (Sec): *The elapsed time in seconds since the routing switch learned the information for the current instance of the indicated multicast flow.*

Expiry Time (Sec): *Indicates the remaining time in seconds before the routing switch ages-out the current flow (group membership). This value decrements until:*

- *Reset by a state refresh packet originating from the upstream multicast router. (The upstream multicast router issues state refresh packets for the current group as long as it either continues to receive traffic for the current flow or receives state refresh packets for the current flow from another upstream multicast router.)*
- *Reset by a new flow for the current multicast group on the VLAN.*
- *The timer expires (reaches **0**). In this case the switch has not received either a state refresh packet or new traffic for the current multicast group, and ages-out (drops) the group entry.*

Multicast Routing Protocol: *Identifies the multicast routing protocol through which the current flow was learned.*

Unicast Routing Protocol: *Identifies the routing protocol through which the routing switch learned the upstream interface for the current multicast flow. The listed protocol will be either **RIP**, **OSPF**, or **Static Route**.*

Downstream Interfaces:

VLAN: *Lists the VID of the VLAN that the routing switch is using to send the outbound packets of the current multicast flow to the next-hop router.*

State: *Indicates whether the outbound VLAN and next-hop router for the current multicast flow are receiving datagrams.*

- **Pruned:** *The routing switch has not detected any joins from the current multicast flow and is not currently forwarding datagrams in the current VLAN.*
- **Forwarding:** *The routing switch has received a join for the current multicast flow and is forwarding datagrams in the current VLAN.*

Up Time (Sec): *Indicates the elapsed time in seconds since the routing switch learned the displayed information about the current multicast flow.*

Expiry Time: *Shows the remaining time in seconds until the Next-Hop routing switch ages-out the current flow (group membership) on the indicated VLAN. Includes the date calculated for the age-out event. This value decrements until:*

- *Reset by a state refresh packet originating from the upstream multicast router. (The upstream multicast router issues state refresh packets for the current group as long as it either continues to receive traffic for the current flow or receives state refresh packets for the current flow from another upstream multicast router.*
- *Reset by a new flow for the current multicast group on the VLAN.*
- *The timer expires (reaches **0**). In this case the switch has not received either a state refresh packet or new traffic for the current multicast group, and ages-out (drops) the group entry.*

Note that the “Next-Hop routing switch” is the next multicast routing switch in the path from the current multicast routing switch to the source for the displayed multicast flow.

```
ProCurve# show ip mroute 239.255.255.5 27.27.30.2

IP Multicast Route Entry

Group Address : 239.255.255.5
Source Address : 27.27.30.2
Source Mask : 255.255.248.0
Neighbor : 13.30.229.1
VLAN : 27
Up Time (sec) : 408
Expiry Time (sec) : 150

Multicast Routing Protocol : PIM-DM
Unicast Routing Protocol : rip

Downstream Interfaces

VLAN State      Up Time (sec)      Expiry Time (sec)
-----
28 pruned       408                98
```

A blank **Neighbor** field indicates that the multicast server is directly connected to the routing switch.

Figure 5-9. Example Output for “5300XL #1” Routing Switch in Figure 5-4 on Page 5-20

Displaying PIM Status

Syntax: show ip pim

Displays PIM status and global parameters.

PIM Status: Shows either **enabled** or **disabled**.

State Refresh Interval (sec): A PIM routing switch originates state refresh messages to inform its neighbors of the active flows it is currently routing. This updates the current flow data on PIM routers that join or rejoin a multicast network after the initial flood and prune. This enables hosts on such routers to join a multicast group without having to wait for a “flood and prune” cycle. PIM routers having the state refresh capability can eliminate all but an initial flood and prune cycle. PIM routers without this capability periodically trigger a flood and prune cycle on the path between the PIM router and the multicast source. (Range: 10 - 300 seconds; Default: 60 seconds.)

Traps: Enables the following SNMP traps:

- **neighbor-loss:** Sends a trap if a neighbor router is lost.
- **hardware-mrt-full:** Sends a trap if the hardware multicast router (MRT) table is full (511 active flows).
- **software-mrt-full:** Sends a trap if the software multicast router (MRT) table is full (511 active flows). This can occur only if the hardware MRT is also full.
- **all:** Enables all of the above traps.

PIM-DM (Dense Mode) on the 5300xl Switches

Displaying PIM Data and Configuration Settings on the Series 5300xl Switches

```
ProCurve# show ip pim

PIM Global Parameters

PIM Status           : enabled
State Refresh Interval (sec) : 45
Traps                 : hardware-mrt-full
```

Figure 5-10. Example Output for the “5304XL #1” Routing Switch in Figure 5-4 on Page 5-20

Syntax: show ip pim [interface]

Lists the PIM interfaces (VLANs) currently configured in the routing switch.

VLAN: *Lists the VID of each VLAN configured on the switch to support PIM-DM.*

IP Address: *Lists the IP addresses of the PIM interfaces (VLANs).*

Mode: *Shows dense only.*

```
ProCurve# show ip pim interface

PIM Interfaces

VLAN IP Address      Mode
-----
25   25.38.10.1        dense
27   27.27.30.1        dense
29   29.29.30.1        dense
```

Figure 5-11. Example Output for the “5304xl #1” Routing Switch in Figure 5-4 on Page 5-20

Syntax: show ip pim [interface [< vid >]]

Displays the current configuration for the specified VLAN (PIM interface). Refer to table 5-1, below.

```

ProCurve# show ip pim interface 29

PIM Interface

VLAN       : 29
IP Address : 29.29.30.1
Mode       : dense

Designated Router :

Hello Interval (sec) : 30
Hello Delay (sec)   : 5

Graft Retry Interval(sec) : 3
Max Graft Retries      : 2
Override Interval (msec) : 2500
Propagation Delay (msec) : 500
SR TTL Threshold       : 0

Lan Prune Delay           : Yes
Lan Delay Enabled        : Yes
State Refresh Capable    : Yes
    
```

Figure 5-12. Example Output for the “5304xl #1” Routing Switch in Figure 5-4 on Page 5-20

Table 5-1. PIM Interface Configuration Settings

Field	Default	Control Command
VLAN	n/a	vlan < vid > ip pim
IP	n/a	vlan < vid > ip pim < all ip-addr >
Mode	dense	n/a; PIM Dense only
Hello Interval (sec)	30	ip pim hello interval < 5 - 30 >
Hello Hold Time	105	The routing switch computes this value from the current “Hello Interval” and includes it in the “Hello” packets the routing switch sends to neighbor routers. Neighbor routers use this value to determine how long to wait for another Hello packet from the routing switch. Refer to the description of the Hello Interval on page 5-15.
Hello Delay	5	vlan < vid > ip pim hello delay < 0 - 5 >
Graft Retry Interval (sec)	3	vlan < vid > ip pim graft-retry-interval < 1 - 10 >

PIM-DM (Dense Mode) on the 5300xl Switches

Displaying PIM Data and Configuration Settings on the Series 5300xl Switches

Field	Default	Control Command
Max Graft Retries	2	vlan < vid > ip pim graft-retries < 1 - 10 >
Override Interval (msec)	2500	vlan < vid > ip pim override-interval < 500 - 6000 >
Propagation Delay (msec)	500	vlan < vid > ip pim propagation-delay < 250-2000 >
SR TTL Threshold (router hops)	0	vlan < vid > ip pim ttl-threshold < 0 - 255 >
LAN Prune Delay	Yes	vlan < vid > ip pim lan-prune-delay
LAN Delay Enabled	No	Shows Yes if all multicast routers on the current VLAN interface enabled LAN-prune-delay. Otherwise shows No .
State Refresh Capable	n/a	Indicates whether the VLAN responds to state refresh packets. The VLAN connected to the multicast source does not receive state refresh packets and thus is not state-refresh capable. Downstream VLANs in Switch Series 5300XL devices are state-refresh capable.

Syntax: show ip pim [mroute]

Shows PIM-specific information from the IP multicast routing table (IP MRT). When invoked without parameters, lists all PIM entries currently in the routing switch's IP MRT.

Group Address: *Lists the multicast group addresses currently active on the routing switch.*

Source Address: *Lists the multicast source address for each Group Address.*

Metric: *Indicates the path cost upstream to the multicast source. Used when multiple multicast routers contend to determine the best path to the multicast source. The lower the value, the better the path. This value is set to 0 (zero) for directly connected routes.*

Metric Pref: *Used when multiple multicast routers contend to determine the path to the multicast source. When this value differs between routers, PIM selects the router with the lowest value. If Metric Pref is the same between contending multicast routers, then PIM selects the router with the lowest Metric value to provide the path for the specified multicast traffic. This value is set to 0 (zero) for directly connected routes.*

(Metric Pref is based on the routing protocol in use: RIP, OSPF, or static routing. Also, different vendors may assign different values for this setting.)

This output shows the routing switch is receiving two multicast groups from an upstream device at 27.27.30.2. The "0" metric shows that the routing switch is directly connected to the multicast source.

```
HPswitch# show ip pim mroute
```

Group Address	Source Address	Metric	Metric Pref
239.255.255.1	27.27.30.2	(0)	0
239.255.255.5	27.27.30.2	(0)	0




Figure 5-13. Example Showing a Routing Switch Detecting two Multicast Groups from a Directly Connected Multicast Server

Syntax: show ip pim [mroute [< multicast-group-address >
< multicast-source-address >]]

Displays the PIM route entry information for the specified multicast group (flow):

Group Address: Lists the specified multicast group address.

Source Address: Lists the specified multicast source address.

Source Mask: Lists the network mask for the multicast source address.

Metric: Lists the number of multicast router hops to the source address.

Metric: Indicates the path cost upstream to the multicast source. Used when multiple multicast routers contend to determine the best path to the multicast source. The lower the value, the better the path.

Metric Pref: Used when multiple multicast routers contend to determine the path to the multicast source. When this value differs between routers, PIM selects the router with the lowest value. If Metric Pref is the same between contending multicast routers, then PIM selects the router with the lowest **Metric** value to provide the path for the specified multicast traffic. (Different vendors assign differing values for this setting.)

Assert Timer: The time remaining until the routing switch ceases to wait for a response from another multicast router to negotiate the best path back to the multicast source. If this timer expires without a response from any contending multicast routers, then the routing switch assumes it is the best path, and the specified multicast group traffic will flow through the routing switch.

DownStream Interfaces:

- **VLAN:** Lists the VID of the destination VLAN on the next-hop multicast router.
- **Prune Reason:** *Identifies the reason for pruning the flow to the indicated VLAN:*
 - **Prune:** *A neighbor multicast router has sent a prune request.*
 - **Assert:** *Another multicast router connected to the same VLAN has been elected to provide the path for the specified multicast group traffic.*
 - **Other:** *Used where the VLAN is in the pruned state for any reason other than the above two reasons (such as no neighbors exist and no directly connected hosts have done joins).*

```
ProCurve# show ip pim mroute 239.255.255.1 27.27.30.2

PIM Route Entry

Group Address   : 239.255.255.1
Source Address  : 27.27.30.2
Source Mask     : 255.255.248.0

Metric         : 3
Metric Pref    : 120
Assert Timer   : 0

DownStream Interfaces

VLAN Prune Reason
-----
28   prune
```

This example displays the MRT data on the first of the two multicast groups shown in figure 5-13 on page 5-31.

Figure 5-14. Example From the “5304xl #1” Routing Switch in Figure 5-4 on Page 5-20 Showing a Multicast Group from a Directly Connected Source

Syntax: show ip pim [neighbor]

Lists PIM neighbor information for all PIM neighbors connected to the routing switch:

IP Address: *Lists the IP address of a neighbor multicast router.*

VLAN: *Lists the VLAN through which the routing switch connects to the indicated neighbor.*

Up Time: *Shows the elapsed time during which the neighbor has maintained a PIM route to the routing switch.*

Expiry Time: *Indicates how long before the routing switch ages-out the current flow (group membership). This value decrements until:*

- *Reset by a state refresh packet originating from the upstream multicast router. (The upstream multicast router issues state refresh packets for the current group as long as it either continues to receive traffic for the current flow or receives state refresh packets for the current flow from another upstream multicast router.*
- *Reset by a new flow for the current multicast group on the VLAN.*

The timer expires (reaches 0). In this case the switch has not received either a state refresh packet or new traffic for the current multicast group, and ages-out (drops) the group entry.

If the IP-ADDR is specified then detailed information for the specified neighbor is shown.

This example simulates output from the "5304XL #1" Routing Switch in Figure 5-4 on Page 5-20. The data identifies the first downstream neighbor ("5300XL #2").

```
ProCurve# show ip pim neighbor
```

PIM Neighbors

IP Address	VLAN	Up Time (sec)	Expiry Time (sec)
29.29.30.2	29	196	89

Figure 5-15. Example of PIM Neighbor Output

Syntax: show ip pim [neighbor [< ip-address >]]

*Lists the same information as **show ip pim neighbor** (page 5-33) for the specified PIM neighbor:*

```
This example simulates output from the "5304XL #1" Routing Switch in Figure 5-4 on Page 5-20. The data is from the first downstream neighbor ("5300XL #2").

ProCurve# show ip pim neighbor 29.29.30.2

PIM Neighbor

  IP Address   : 29.29.30.2
  VLAN        : 29

  Up Time (sec)      : 26
  Expiry Time (sec) : 79
```

Figure 5-16. Example From the "5304xl #1" Routing Switch in Figure 5-4 on Page 5-20 Showing a Specific Neighbor ("5300xl #2")

Operating Notes

PIM Routers without State Refresh Messaging Capability. A PIM router without a state refresh messaging capability learns of currently active flows in a multicast network through periodic flood and prune cycles on the path back to the source. The Switch Series 5300XL devices sense downstream multicast routers that do not have the state refresh capability and will periodically flood active multicast groups to these devices. This periodic flooding is not necessary if all of the downstream multicast routers are ProCurve 5300XL devices. (The ProCurve Routing Switch Series 9300 and the routers offered by some other vendors do not offer the state refresh capability.)

Flow Capacity. The routing switch provides an ample multicast environment, supporting 1022 multicast flows in hardware across a maximum of 64 VLANs. (A flow comprises a unicast source address and a multicast group address, regardless of the number of active hosts belonging to the multicast group at any given time.) While the typical multicast environment should not normally exceed 1022 flows, the routing switch can support up to 978 additional flows in software, depending on available system resources. (Because the switch processes flows in hardware much faster than in software, you may notice slower processing times for flows occurring in software.) Also, while the routing switch can support up to 2,000 flows, the total demand on system resources from the combined use of more than 1,022 simultaneous flows, a high number of VLANs supporting multicast routing, and/or other, resource-intensive features can oversubscribe memory resources, which reduces the number of flows the routing switch can support in software. That is, the switch does not route flows in software that oversubscribe current memory resources. If the routing switch regularly exceeds the hardware limit of 1022 flows and begins routing flows in software, you may want to move some hosts that create multicast demand to another routing switch, or reduce the number of VLANs on the routing switch by moving some VLANs to another routing switch. Note that the routing switch generates a log message if it either routes a flow in software or drops a flow intended for software routing because memory is oversubscribed. (Refer to “Messages Related to PIM Operation” on page 5-37.)

IGMP Traffic High-Priority Disabled. Enabling IP multicast routing to support PIM-DM operation has the effect of disabling IGMP traffic high-priority, if configured. (Refer to “Configuring IGMP Traffic Priority” on page 4-10.)

ACLs and PIM. The switch allows ACL filtering on unicast addresses, but not on multicast addresses. Also, an ACL does not take effect on a flow if the flow began before the ACL was configured.

When To Enable IGMP on a VLAN. When PIM is enabled on a VLAN, it is not necessary to also enable IGMP unless there may be Joins occurring on that VLAN. But if IGMP is enabled on a VLAN, you must also enable PIM if you want that VLAN to participate in multicast routing.

IP Address Removed. If you remove the IP address for a VLAN, the switch automatically removes the PIM configuration for that VLAN.

Troubleshooting

Symptom: Noticeable slowdown in some multicast traffic. If the switch is supporting more than 1022 active flows. This generates the message `Unable to learn HW IP multicast groups, table FULL` in the Event Log because there is no room in the hardware Multicast Routing Table to add another Multicast Group. Software will route any multicast packets sent to multicast groups that are not in the hardware Multicast Routing Table, but it will be slower and packets may be dropped if the data rate is greater than 3000 packets per second. Refer to “Flow Capacity” on page 5-35.

Note that the PIM protocol uses one MRT entry for every IP multicast source/group pair that it is routing. An entry is not used if the multicast flow is bridged and not routed. Entries in this table are automatically aged out if they are unused for a period of time.

Heavy Memory Usage. Heavy use of PIM (many S/G flows over many VLANs) combined with other memory-intensive features, can oversubscribe memory resources and impact overall performance. If available memory is exceeded, the switch drops any new multicast flows and generates appropriate Event Log messages. Corrective actions can include reducing the number of VLANs on the 5300xl switch by moving some VLANs to another device, free up system resources by disabling another, non-PIM feature, and/or moving some hosts to another device. For more information, refer to “Operating Notes” on page 5-34 and “Messages Related to PIM Operation” on page 5-37.

IPv4 Table Operation. The IPv4 table, which contains the active IP multicast addresses the switch is currently supporting, has 128k entries. However, the IPv4 table also contains IP host entries for every IP source or destination that the switch has learned, as well as ACL flow entries. Entries in this table are generally aged out if they are unused for 5 minutes or more.

Messages Related to PIM Operation

These messages appear in the Event Log and, if Syslog Debug is configured, in the designated Debug destinations.

Note

The <counter> value displayed at the end of each PIM Event Log message (and SNMP trap messages, if trap receivers are configured) indicates the number of times the switch has detected a recurring event since the last reboot. For more information, refer to “Using the Event Log To Identify Problem Sources” in the “Troubleshooting” appendix of the February, 2004 (or later) version of the *Management and Configuration Guide* for your switch.

Message	Meaning
<alpha-string> pkt, src IP<ip-addr> vid <vlan-id> (not a nbr) (<counter>)	A PIM packet arrived from another router for which no neighbor was found. May indicate a misconfiguration between the sending and receiving router. May also occur if a connected router is disconnected, then reconnected.
Bad TTL in State Refresh pkt from IP <source-ip-addr> (<counter>)	The switch detected a TTL of 0 (zero) in the PIM portion of a state refresh packet. (Note that this is not the IP TTL.)
Failed alloc of HW <alpha-str> for flow <multicast-address>, <source-address> (<dup-msg-cnt>)	There are more than 1022 active flows. The switch routes the excess through software, which processes traffic at a slower rate. If this will be an ongoing or chronic condition, transfer some of the flows to another router.
Failed to alloc a PIM <data-type> pkt (<counter>)	The router was unable to allocate memory for a PIM control packet. Router memory is oversubscribed. Reduce the number of VLANs or increase the hello delay and/or the override interval to reduce the number of simultaneous packet transmissions. Note that if the number of flows exceeds 1022, the excess flows are routed in software, which reduces the number of packet transmissions. In this case, reducing the number of flows by moving some clients to other routers can help.
Failed to initialize <text-str> as a call back routine (<counter>)	Indicates an internal error. Report the incident to your HP customer care center and re-install the router software.
I/F configured with IP <ip-address> on vid <vlan-id> (<counter>)	Indicates that the interface (VLAN) has been configured with the indicated IP address. At boot-up or when an IP address is changed, the switch generates this message for each PIM-configured VLAN.

PIM-DM (Dense Mode) on the 5300xl Switches
 Messages Related to PIM Operation

Message	Meaning
I/F removal with IP < <i>ip-addr</i> > on vid < <i>vlan-id</i> > (< <i>counter</i> >)	Indicates that a PIM interface (VLAN) has been removed from the router as a result of an IP address change or removal.
MCAST flow < <i>multicast-address</i> > < <i>source-address</i> > not rteing (rsc low) (< <i>counter</i> >)	The indicated multicast flow is not routing. The routing switch is low on memory resources as a result of too many flows for the number of configured VLANs. Remedies include one or more of the following: <ul style="list-style-type: none"> • Reduce the number of configured VLANs by moving some VLANs to another router. • Free up system resources by disabling another feature, such as one of the spanning-tree protocols or either the RIP or the OSPF routing protocol. (Unless you are using static routes, you will need to retain a minimum of one unicast routing protocol.) Another option that may help is to reduce the number of configured QoS filters. • Move some hosts that create multicast demand to another router.
MCAST MAC add for < <i>mac-address</i> > failed (< <i>counter</i> >)	Indicates a hardware problem. Check the cabling and router ports.
Multicast Hardware Failed to Initialize (< <i>counter</i> >)	Indicates a hardware failure that halts hardware processing of PIM traffic. The software will continue to process PIM traffic at a slower rate. Contact your HP customer care center.
No IP address configured on VID < <i>vlan-id</i> > (< <i>dup-msg-cnt</i> >)	PIM has detected a VLAN without an IP address. Configure an IP address on the indicated VLAN.
Pkt dropped from < <i>ip-address</i> >, (< <i>cause</i> >) vid < <i>vlan-id</i> > (< <i>counter</i> >)	A PIM packet from < <i>ip-address</i> > was dropped due to one of the following causes: <ul style="list-style-type: none"> • No PIM interface on the VLAN • Bad packet length • Bad IP header length • Bad IP total length
Pkt rcvd with a cksum error from < <i>ip-addr</i> > (< <i>counter</i> >)	A packet having a checksum error was received from < <i>ip-address</i> >. Check the cabling and ports on the local and the remote routers.
Rcvd incorrect hello from < <i>ip-addr</i> > (< <i>counter</i> >)	Indicates receipt of a malformed hello packet. (That is, the packet does not match the current specification.) Ensure that compatible versions of PIM-DM are being used.
Rcvd < <i>text-str</i> > pkt with bad len from < <i>ip-addr</i> > (< <i>counter</i> >)	A peer router may be sending incorrectly formatted PIM packets.
Rcvd hello from < <i>ip-address</i> > on vid < <i>vlan-id</i> > (< <i>counter</i> >)	Indicates a misconfiguration where two routers are directly connected with different subnets on the same connected interface.

Message	Meaning
Rcvd pkt from rtr < <i>ip-address</i> >, unkwn pkt type < <i>value</i> > (< <i>counter</i> >)	A packet received from the router at < <i>ip-address</i> > is an unknown PIM packet type. (The < <i>value</i> > variable is the numeric value received in the packet.)
Rcvd pkt ver# < <i>ver-num</i> >, from < <i>ip-address</i> >, expected < <i>ver-num</i> > (< <i>counter</i> >)	The versions of PIM-DM on the sending and receiving routers do not match. Differing versions will typically be compatible, but features not supported in both versions will not be available.
Rcvd unkwn addr fmly < <i>addr-type</i> > in < <i>text-str</i> > pkt from < <i>ip-addr</i> > (< <i>counter</i> >)	The router received a PIM packet with an unrecognized encoding. As of February, 2004, the router recognizes IPv4 encoding.
Rcvd unkwn opt < <i>opt-nbr</i> > in < <i>text-string</i> > pkt from < <i>ip-addr</i> > (< <i>counter</i> >)	The router received a PIM packet carrying an unknown PIM option. The packet may have been generated by a newer version of PIM-DM, or is corrupt. In most cases, normal PIM-DM operation will continue.
Send error(< <i>failure-type</i> >) on < <i>packet-type</i> > pkt on VID < <i>vid</i> > (< <i>counter</i> >)	Indicates a send error on a packet. This can occur if a VLAN went down right after the packet was sent. The message indicates the failure type, the packet type, and the VLAN ID on which the packet was sent.
Unable to alloc < <i>text-str</i> > table (< <i>counter</i> >)	The router was not able to create some tables PIM-DM uses. Indicates that the router is low on memory resources. Remedies include one or more of the following: <ul style="list-style-type: none"> • Reduce the number of configured VLANs by moving some VLANs to another router. • Free up system resources by disabling another feature, such as one of the spanning-tree protocols or either the RIP or the OSPF routing protocol. (Unless you are using static routes, you will need to retain a minimum of one unicast routing protocol.) Another option that may help is to reduce the number of configured QoS filters. • Move some hosts that create multicast demand to another router.
Unable to alloc a buf of size < <i>bytes</i> > for < <i>data-flow</i> > (< <i>counter</i> >)	Multicast routing is unable to acquire memory for a flow. Router memory is oversubscribed. Reduce the number of VLANs or the number of features in use. Remedies include one or more of the following: <ul style="list-style-type: none"> • Reduce the number of configured VLANs by moving some VLANs to another router. • Free up system resources by disabling another feature, such as one of the spanning-tree protocols or either the RIP or the OSPF routing protocol. (Unless you are using static routes, you will need to retain a minimum of one unicast routing protocol.) Another option that may help is to reduce the number of configured QoS filters. • Move some hosts that create multicast demand to another router.

Message	Meaning
Unable to alloc a msg buffer for <text-message> (<counter>)	Multicast routing is unable to acquire memory for a flow. Router memory is oversubscribed. Reduce the number of VLANs or the number of features in use. Remedies include one or more of the following: <ul style="list-style-type: none">• Reduce the number of configured VLANs by moving some VLANs to another router.• Free up system resources by disabling another feature, such as one of the spanning-tree protocols or either the RIP or the OSPF routing protocol. (Unless you are using static routes, you will need to retain a minimum of one unicast routing protocol.) Another option that may help is to reduce the number of configured QoS filters.• Move some hosts that create multicast demand to another router.

Applicable RFCs

PIM on the Switch Series 5300XL devices is compatible with these RFCs:

- RFC 3376 - Internet Group Management Protocol, Version 3
- RFC 2365 - Administratively Scoped IP Multicast
- RFC 2932 - Multicast Routing MIB, *with exceptions (Refer to "Exceptions to Support for RFC 2932 - Multicast Routing MIB".)*
- RFC 2933 - IGMP MIB
- RFC 2934 - Protocol Independent Multicast MIB for IPv4
- draft-ietf-ssm-arch-01.txt - Source-Specific Multicast for IP (draft specification, expires May 2003)

Exceptions to Support for RFC 2932 - Multicast Routing MIB

These MIB objects are not supported in the 5300XL routing switch.

ipMRouteInterfaceRateLimit
ipMRouteInterfaceInMcastOctets
ipMRouteInterfaceOutMcastOctets
ipMRouteInterfaceHCInMcastOctets
ipMRouteInterfaceHCOutMcastOctets
ipMRouteBoundaryTable
ipMRouteBoundaryEntry
ipMRouteBoundaryIfIndex
ipMRouteBoundaryAddress
ipMRouteBoundaryAddressMask
ipMRouteBoundaryStatus OBJECT-TYPE
ipMRouteScopeNameTable
ipMRouteScopeNameEntry
ipMRouteScopeNameAddress
ipMRouteScopeNameAddressMask
ipMRouteScopeNameLanguage
ipMRouteScopeNameString
ipMRouteScopeNameDefault
ipMRouteScopeNameStatus

— *This page is intentionally unused.* —