2

Multimedia Traffic Control with IP Multicast (IGMP)

Contents

Overview

This chapter describes multimedia traffic control with IP multicast (IGMP) to reduce unnecessary bandwidth usage on a per-port basis, and how to configure it with the switch's built-in interfaces:

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 ProCurve Web Browser Interface
- Chapter 6, "Switch Memory and Configuration"

Note

The use of static multicast filters is described in the chapter titled "Traffic/ Security Filters" in the *Access Security Guide* for your ProCurve switch.

IGMP General Operation and Features

IGMP Features

Feature	Default	Menu	CLI
view igmp configuration	n/a	—	page 2-7
show igmp status for multicast groups used by the selected VLAN	n/a	—	Yes
enabling or disabling IGMP (Requires VLAN ID Context)	disabled	—	page 2-9
per-port packet control	auto	—	page 2-10
IGMP traffic priority	normal	—	page 2-11
querier	enabled	—	page 2-11
fast-leave	disabled	—	page 2-14

In a network where IP multicast traffic is transmitted for various multimedia applications, you can use the switch to reduce unnecessary bandwidth usage on a per-port basis by configuring IGMP (Internet Group Management Protocol controls). In the factory default state (IGMP disabled), the switch simply floods all IP multicast traffic it receives on a given VLAN through all ports on that VLAN (except the port on which it received the traffic). This can result in significant and unnecessary bandwidth usage in networks where IP multicast traffic is a factor. Enabling IGMP allows the ports to detect IGMP queries and report packets and manage IP multicast traffic through the switch.

IGMP is useful in multimedia applications such as LAN TV, desktop conferencing, and collaborative computing, where there is multipoint communication; that is, communication from one to many hosts, or communication originating from many hosts and destined for many other hosts. In such multipoint applications, IGMP will be configured on the hosts, and multicast traffic will be generated by one or more servers (inside or outside of the local network). Switches in the network (that support IGMP) can then be configured to direct the multicast traffic to only the ports where needed. If multiple VLANs are configured, you can configure IGMP on a per-VLAN basis.

Enabling IGMP allows detection of IGMP queries and report packets in order to manage IP multicast traffic through the switch. If no other querier is detected, the switch will then also function as the querier. (If you need to disable the querier feature, you can do so through the IGMP configuration MIB. Refer to “Changing the Querier Configuration Setting” on page 2-11.)

Note

IGMP configuration on the switches covered in this guide operates at the VLAN context level. If you are not using VLANs, then configure IGMP in VLAN 1 (the default VLAN) context.

IGMP Terms

- **IGMP Device:** A switch or router running IGMP traffic control features.
- **IGMP Host:** An end-node device running an IGMP (multipoint, or multicast communication) application.
- **Querier:** A required IGMP device that facilitates the IGMP protocol and traffic flow on a given LAN. This device tracks which ports are connected to devices (IGMP clients) that belong to specific multicast groups, and triggers updates of this information. A querier uses data received from the queries to determine whether to forward or block multicast traffic on specific ports. When the switch has an IP address on a given VLAN, it automatically operates as a Querier for that VLAN if it does not detect a multicast router or another switch functioning as a Querier. When enabled (the default state), the switch's querier function eliminates the need for a multicast router. In most cases, ProCurve recommends that you leave this parameter in the default "enabled" state even if you have a multicast router performing the querier function in your multicast group. For more information, see "How IGMP Operates" on page 2-12.

IGMP Operating Features

Basic Operation

In the factory default configuration, IGMP is disabled. To enable IGMP

- If multiple VLANs are not configured, you configure IGMP on the default VLAN (DEFAULT_VLAN; VID = 1).
- If multiple VLANs are configured, you configure IGMP on a per-VLAN basis for every VLAN where this feature is to be used.

Enhancements

With the CLI, you can configure these additional options:

- **Forward with High Priority.** Disabling this parameter (the default) causes the switch or VLAN to process IP multicast traffic, along with other traffic, in the order received (usually, normal priority). Enabling this parameter causes the switch or VLAN to give a higher priority to IP multicast traffic than to other traffic.
- **Auto/Blocked/Forward:** You can use the console to configure individual ports to any of the following states:
  - **Auto** (the default): Causes the switch to interpret IGMP packets and to filter IP multicast traffic based on the IGMP packet information for ports belonging to a multicast group. This means that IGMP traffic will be forwarded on a specific port only if an IGMP host or multicast router is connected to the port.
  - **Blocked:** Causes the switch to drop all IGMP transmissions received from a specific port and to block all outgoing IP Multicast packets for that port. This has the effect of preventing IGMP traffic from moving through specific ports.
  - **Forward:** Causes the switch to forward all IGMP and IP multicast transmissions through the port.
- **Operation With or Without IP Addressing:** This feature helps to conserve IP addresses by enabling IGMP to run on VLANs that do not have an IP address. See “Operation With or Without IP Addressing” on page 2-13.
- **Querier Capability:** The switch performs this function for IGMP on VLANs having an IP address when there is no other device in the VLAN acting as querier. See “Using the Switch as Querier” on page 2-26.

Notes

Whenever IGMP is enabled, the switch generates an Event Log message indicating whether querier functionality is enabled.

IP multicast traffic groups are identified by IP addresses in the range of 224.0.0.0 to 239.255.255.255. Also, incoming IGMP packets intended for reserved, or "well-known" multicast addresses automatically flood through all ports (except the port on which the packets entered the switch). For more on this topic, see "Excluding Well-Known or Reserved Multicast Addresses from IP Multicast Filtering" on page 2-27.

For more information, refer to "How IGMP Operates" on page 2-12.

Number of IP Multicast Addresses Allowed

The total of IGMP filters (addresses) and static multicast filters together is 2,047 (if data driven) or 2,048 otherwise, depending on the current **max-vlans** configuration. If multiple VLANs are configured, then each filter is counted once per VLAN in which it is used.

Number of Multicast Filters Allowed

The number of multicast filters allowed depends on the number of configured VLANS:

- 16 multicast filters if VLANs <= 1024
- 8 multicast filters if VLANs > 1024

CLI: Configuring and Displaying IGMP

IGMP Commands Used in This Section

show ip igmp configuration	page 2-7
ip igmp	page 2-9
high-priority-forward	page 2-11
auto <[ethernet] *<port-list>*	page 2-10
blocked <[ethernet] *<port-list>*	page 2-10
forward <[ethernet] *<port-list>*	page 2-10
querier	page 2-11
show ip igmp	Refer to the section titled "Internet Group Management Protocol (IGMP) Status" in appendix B of the *Management and Configuration Guide* for your switch.

Viewing the Current IGMP Configuration. This command lists the IGMP configuration for all VLANs configured on the switch or for a specific VLAN.

Syntax: show ip igmp config

Displays IGMP configuration for all VLANs on the switch.

show ip igmp vlan *< vid >* config

Displays IGMP configuration for a specific VLAN on the switch, including per-port data.

(For IGMP operating status, refer to the section titled "Internet Group Management Protocol (IGMP) Status" in appendix B, "Monitoring and Analyzing Switch Operation" of the *Management and Configuration Guide* for you switch.)

For example, suppose you have the following VLAN and IGMP configurations on the switch:

VLAN ID	VLAN Name	IGMP Enabled	Forward with High Priority	Querier
1	DEFAULT_VLAN	Yes	No	No
22	VLAN-2	Yes	Yes	Yes
33	VLAN-3	No	No	No

You could use the CLI to display this data as follows:

```
ProCurve> show ip igmp config
 IGMP Service
  VLAN ID      VLAN NAME    IGMP Enabled Forward with High Priority Querier
  ------------ ------------ ------------ -------------------------- -------
  1            DEFAULT_VLAN Yes          No                         No
  22           VLAN-2       Yes          Yes                        Yes
  33           VLAN-3       No           No                         Yes
```

Figure 2-1. Example Listing of IGMP Configuration for All VLANs in the Switch

The following version of the **show ip igmp** command includes the VLAN ID (*vid*) designation, and combines the above data with the IGMP per-port configuration:

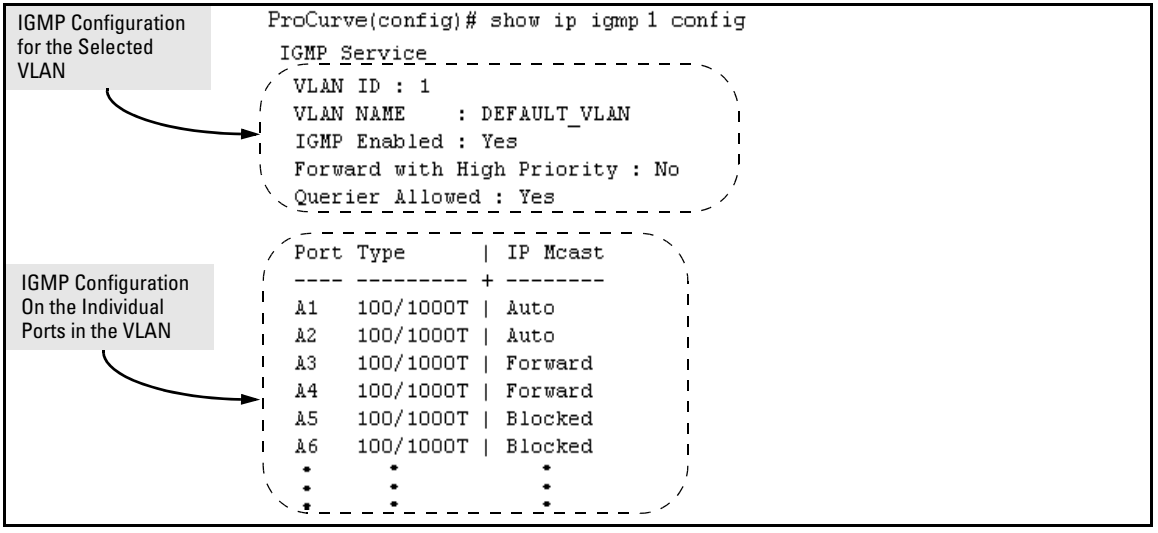

Figure 2-2. Example Listing of IGMP Configuration for A Specific VLAN

Enabling or Disabling IGMP on a VLAN. You can enable IGMP on a VLAN, along with the last-saved or default IGMP configuration (whichever was most recently set), or you can disable IGMP on a selected VLAN.

Syntax: [no] ip igmp

Enables IGMP on a VLAN. Note that this command must be executed in a VLAN context.

For example, here are methods to enable and disable IGMP on the default VLAN (VID = 1).

```
ProCurve(config)# vlan 1 ip igmp
```

Enables IGMP on VLAN 1.

```
ProCurve(vlan-1)# ip igmp
```

Same as above.

```
ProCurve(config)# no vlan 1 ip igmp
```

Disables IGMP on vlan 1.

Note

If you disable IGMP on a VLAN and then later re-enable IGMP on that VLAN, the switch restores the last-saved IGMP configuration for that VLAN. For more on how switch memory operates, refer to the chapter titled "Switch Memory and Configuration" in the *Management and Configuration Guide* for your switch.

You can also combine the ip igmp command with other IGMP-related commands, as described in the following sections.

Configuring Per-Port IGMP Traffic Filters.

Syntax: vlan < *vid* > ip igmp [auto < *port-list* > | blocked < *port-list* > | forward < *port-list* >]

Used in the VLAN context, this command specifies how each port should handle IGMP traffic. (Default: **auto**.*)*

Note: *Where a static multicast filter is configured on a port, and an IGMP filter created by this command applies to the same port, the IGMP filter overrides the static multicast filter for any inbound multicast traffic carrying the same multicast address as is configured in the static filter. (Refer to the section titled "Filter Types and Operation" in the "Port Traffic Controls" chapter of the* Management and Configuration Guide *for your switch.*

For example, suppose you wanted to configure IGMP as follows for VLAN 1 on the 100/1000T ports on a module in slot 1:

Ports A1-A2	auto	Filter multicast traffic. Forward IGMP traffic to hosts on these ports that belong to the multicast group for which the traffic is intended. (Also forward any multicast traffic through any of these ports that is connected to a multicast router.)
Ports A3-A4	forward	Forward all multicast traffic through this port.
Ports A5-A6	blocked	Drop all multicast traffic received from devices on these ports, and prevent any outgoing multicast traffic from moving through these ports.

Depending on the privilege level, you could use one of the following commands to configure IGMP on VLAN 1 with the above settings:

```
ProCurve(config)# vlan 1 ip igmp auto a1,a2 forward a3,a4
blocked a5,a6

ProCurve(config)# ip igmp auto a1,a2 forward a3,a4 blocked
a5,a6
```

The following command displays the VLAN and per-port configuration resulting from the above commands.

```
ProCurve> show igmp vlan 1 config
```

Configuring IGMP Traffic Priority.

Syntax: vlan < *vid* > ip igmp high-priority-forward

This command assigns "high" priority to IGMP traffic or returns a high-priority setting to "normal" priority. (The traffic will be serviced at its inbound priority.) (Default: **normal***.)*

```
ProCurve(config)# vlan 1 ip igmp high-priority-forward
```

Configures high priority for IGMP traffic on VLAN 1.

```
ProCurve(vlan-1)# ip igmp high-priority-forward
```

Same as above command, but in the VLAN 1 context level.

```
ProCurve(vlan 1)# no ip igmp high-priority-forward
```

Returns IGMP traffic to "normal" priority.

```
ProCurve> show ip igmp config
```

Show command to display results of above high-priority commands.

Configuring the Querier Function.

Syntax: [no] vlan <*vid*> ip igmp querier

This command disables or re-enables the ability for the switch to become querier if necessary. The **no** *version of the command disables the querier function on the switch. The* **show ip igmp config** *command displays the current querier command. (Default Querier Capability: Enabled.)*

How IGMP Operates

The Internet Group Management Protocol (IGMP) is an internal protocol of the Internet Protocol (IP) suite. IP manages multicast traffic by using switches, multicast routers, and hosts that support IGMP. (In Hewlett-Packard's implementation of IGMP, a multicast router is not necessary as long as a switch is configured to support IGMP with the **querier** feature enabled.) A set of hosts, routers, and/or switches that send or receive multicast data streams to or from the same source(s) is termed a *multicast group*, and all devices in the group use the same multicast group address. The multicast group running version 2 of IGMP uses three fundamental types of messages to communicate:

- **Query:** A message sent from the querier (multicast router or switch) asking for a response from each host belonging to the multicast group. If a multicast router supporting IGMP is not present, then the switch must assume this function in order to elicit group membership information from the hosts on the network. (If you need to disable the querier feature, you can do so through the CLI, using the IGMP configuration MIB. See "Configuring the Querier Function" on page 2-11.)
- **Report (Join):** A message sent by a host to the querier to indicate that the host wants to be or is a member of a given group indicated in the report message.
- **Leave Group:** A message sent by a host to the querier to indicate that the host has ceased to be a member of a specific multicast group.

Note on IGMP version 3 support

When an IGMPv3 Join is received by the switch, it accepts the host request and begins to forward the IGMP traffic. This means that ports which have not joined the group and are not connected to routers or the IGMP Querier will not receive the group's multicast traffic.

The switch does not support the IGMPv3 "Exclude Source" or "Include Source" options in the Join Reports. Rather, the group is simply joined from all sources.

The switch does not support becoming a version 3 Querier. It will become a version 2 Querier in the absence of any other Querier on the network.

An IP multicast packet includes the multicast group (address) to which the packet belongs. When an IGMP client connected to a switch port needs to receive multicast traffic from a specific group, it joins the group by sending an IGMP report (join request) to the network. (The multicast group specified

in the join request is determined by the requesting application running on the IGMP client.) When a networking device with IGMP enabled receives the join request for a specific group, it forwards any IP multicast traffic it receives for that group through the port on which the join request was received. When the client is ready to leave the multicast group, it sends a Leave Group message to the network and ceases to be a group member. When the leave request is detected, the appropriate IGMP device will cease transmitting traffic for the designated multicast group through the port on which the leave request was received (as long as there are no other current members of that group on the affected port).

Thus, IGMP identifies members of a multicast group (within a subnet) and allows IGMP-configured hosts (and routers) to join or leave multicast groups.

IGMP Data. To display data showing active group addresses, reports, queries, querier access port, and active group address data (port, type, and access), refer to the section titled "Internet Group Management Protocol (IGMP) Status" in appendix B, "Monitoring and Analyzing Switch Operation" of the *Management and Configuration Guide* for you switch.).

Operation With or Without IP Addressing

You can configure IGMP on VLANs that do not have IP addressing. The benefit of IGMP without IP addressing is a reduction in the number of IP addresses you have to use and configure. This can be significant in a network with a large number of VLANs. The limitation on IGMP without IP addressing is that the switch cannot become Querier on any VLANs for which it has no IP address—so the network administrator must ensure that another IGMP device will act as Querier. It is also advisable to have an additional IGMP device available as a backup Querier. See the following table.

Table 2-1.Comparison of IGMP Operation With and Without IP Addressing

IGMP Function Available With IP Addressing Configured on the VLAN	Available *Without* IP Addressing?	Operating Differences Without an IP Address
Forward multicast group traffic to any port on the VLAN that has received a join request for that multicast group.	Yes	None
Forward join requests (reports) to the Querier.	Yes	None
Configure individual ports in the VLAN to **Auto** (the default)/**Blocked**, or **Forward**.	Yes	None

IGMP Function Available With IP Addressing Configured on the VLAN	Available *Without* IP Addressing?	Operating Differences Without an IP Address
Configure IGMP traffic forwarding to normal or high-priority forwarding.	Yes	None
Age-Out IGMP group addresses when the last IGMP client on a port in the VLAN leaves the group.	Yes	Requires that another IGMP device in the VLAN has an IP address and can operate as Querier. This can be a multicast router or another switch configured for IGMP operation. (ProCurve recommends that the VLAN also include a device operating as a backup Querier in case the device operating as the primary Querier fails for any reason.
Support Fast-Leave IGMP and Forced Fast-Leave IGMP (below).	Yes	
Support automatic Querier election.	No	Querier operation not available.
Operate as the Querier.	No	Querier operation not available.
Available as a backup Querier.	No	Querier operation not available.

Automatic Fast-Leave IGMP

Fast-Leave IGMP. Depending on the switch model, Fast-Leave is enabled or disabled in the default configuration.

Switch Model or Series	Data-Driven IGMP Included?	IGMP Fast-Leave Setting	Default IGMP Behavior
Switch 8212zl Switch 6400cl Switch 6200yl Switch 5400zl Switch 5300xl Switch 4200vl Switch 3500yl Switch 3400cl Switch 2500	Yes	Always Enabled	Drops unjoined mulitcast traffic except for always-fowarded traffic toward the Querier or multicast routers, and out of IGMP-forward ports. Selectively forwards joined multicast traffic.
Switch 2600 Switch 2600-PWR Switch 4100gl Switch 6108	No	Disabled in the Default Configuration	IGMP Fast-Leave disabled in the default configuration. Floods unjoined multicast traffic to all ports. Selectively forwards joined multicast traffic.

On switches that do not support Data-Driven IGMP, unregistered multicast groups are flooded to the VLAN rather than pruned. In this scenario, Fast-Leave IGMP can actually increase the problem of multicast flooding by removing the IGMP group filter before the Querier has recognized the IGMP

leave. The Querier will continue to transmit the multicast group during this short time, and because the group is no longer registered the switch will then flood the multicast group to all ports.

On ProCurve switches that do support Data-Driven IGMP ("Smart" IGMP), when unregistered multicasts are received the switch automatically filters (drops) them. Thus, the sooner the IGMP Leave is processed, the sooner this multicast traffic stops flowing.

Because of the multicast flooding problem mentioned above, the IGMP Fast-Leave feature is disabled by default on all ProCurve switches that do not support Data-Driven IGMP. (See the table above.) The feature can be enabled on these switches via an SNMP set of this object:

hpSwitchIgmpPortForceLeaveState.<*vid*>.<*port number*>

However, this is not recommended as this will increase the amount of multicast flooding during the period between the client's IGMP Leave and the Querier's processing of that Leave. For more information on this topic refer to "Forced Fast-Leave IGMP" on page page 2-17.

Automatic Fast-Leave Operation. If a switch port has the following characteristics, then the Fast-Leave operation will apply:

1. Connected to only one end node
2. The end node currently belongs to a multicast group; i.e. is an IGMP client
3. The end node subsequently leaves the multicast group

Then the switch does not need to wait for the Querier status update interval, but instead immediately removes the IGMP client from its IGMP table and ceases transmitting IGMP traffic to the client. (If the switch detects multiple end nodes on the port, automatic Fast-Leave does not activate—regardless of whether one or more of these end nodes are IGMP clients.)

In the next figure, automatic Fast-Leave operates on the switch ports for IGMP clients “3A” and “5A”, but not on the switch port for IGMP clients “7A” and 7B, Server “7C”, and printer “7D”.

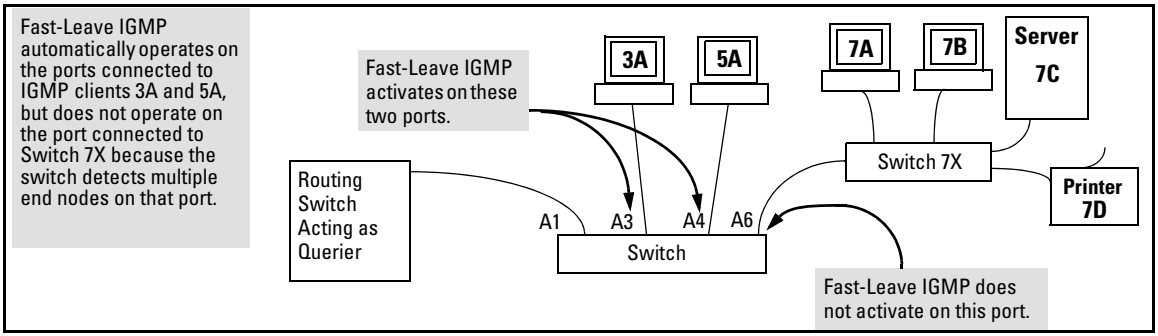

Figure 2-3. Example of Automatic Fast-Leave IGMP Criteria

When client “3A” running IGMP is ready to leave the multicast group, it transmits a Leave Group message. Because the switch knows that there is only one end node on port A3, it removes the client from its IGMP table and halts multicast traffic (for that group) to port A3. If the switch is not the Querier, it does not wait for the actual Querier to verify that there are no other group members on port A3. If the switch itself is the Querier, it does not query port A3 for the presence of other group members.

Note that Fast-Leave operation does not distinguish between end nodes on the same port that belong to different VLANs. Thus, for example, even if all of the devices on port A6 in figure 2-3 belong to different VLANs, Fast-Leave does not operate on port A6.

Default (Enabled) IGMP Operation Solves the “Delayed Leave” Problem. Fast-leave IGMP is enabled by default. When Fast-leave is disabled and multiple IGMP clients are connected to the same port on an IGMP device (switch or router), if only one IGMP client joins a given multicast group, then later sends a Leave Group message and ceases to belong to that group, the switch automatically retains that IGMP client in its IGMP table and continues forwarding IGMP traffic to the IGMP client until the Querier triggers confirmation that no other group members exist on the same port. This delayed leave operation means that the switch continues to transmit unnecessary multicast traffic through the port until the Querier renews multicast group status.

Configuring Fast-Leave IGMP.

Syntax: [no] ip igmp fastleave < *port-list* >

Enables IGMP fast-leaves on the specified ports in the selected VLAN. The **no** *form of the command disables IGMP fast-leave on the specified ports in the selected VLAN. Use* **show running** *to display the ports per-VLAN on which Fast-Leave is disabled.*

Forced Fast-Leave IGMP

When enabled, Forced Fast-Leave IGMP speeds up the process of blocking unnecessary IGMP traffic to a switch port that is connected to multiple end nodes. (This feature does not activate on ports where the switch detects only one end node). For example, in figure 2-3, even if you configured Forced Fast-Leave on all ports in the switch, the feature would activate only on port A6 (which has multiple end nodes) when a Leave Group request arrived on that port.

When a port having multiple end nodes receives a Leave Group request from one end node for a given multicast group "X", Forced Fast-Leave activates and waits a small amount of time to receive a join request from any other group "X" member on that port. If the port does not receive a join request for that group within the forced-leave interval, the switch then blocks any further group "X" traffic to the port.

Configuring Forced Fast-Leave IGMP

Syntax: [no] vlan < *vid* > ip igmp forcedfastleave <port-list>

Enables IGMP Forced Fast-Leave on the specified ports in the selected VLAN, even if they are cascaded. (Default: Disabled.) The **no** *form of the command disables Forced Fast-Leave on the specified ports in the selected VLAN. Use* **show running** *to display the ports per-VLAN on which Forced Fast-Leave is enabled.*

To view a non-default IGMP forced fast-leave configuration on a VLAN, use the **show running-config** command. (The **show running-config** output does not include forced fast-leave if it is set to the default of **0**.)

Forced fast-leave can be used when there are multiple devices attached to a port.

Configuring Delayed Group Flush

When enabled, this feature continues to filter IGMP groups for a specified additional period of time after IGMP leaves have been sent. The delay in flushing the group filter prevents unregistered traffic from being forwarded by the server during the delay period. In practice, this is rarely necessary on the switches covered in this guide, which support data-driven IGMP. (Data-Driven IGMP, which is enabled by default, prunes off any unregistered IGMP streams detected on the switch.)

Syntax: igmp delayed-flush < *time-period* >

Where leaves have been sent for IGMP groups, enables the switch to continue to flush the groups for a specified period of time. This command is applied globally to all IGMP-configured VLANs on the switch. Range: **0 - 255***; Default: Disabled (***0***).*

Syntax: show igmp delayed-flush

Displays the current **igmp delayed-flush** *setting.*

IGMP Proxy Forwarding

Note

For more information about PIM-DM and PIM-SM, see the chapters "*PIM-DM (Dense Mode)*" and "*PIM-SM (Sparse Mode)*" in this guide.

When a network has a border router connecting a PIM-SM domain to a PIM-DM domain, the routers that are completely within the PIM-DM domain have no way to discover multicast flows in the PIM-SM domain. When an IGMP join occurs on a router entirely within the PIM-DM domain for a flow that originates within the PIM-SM domain, it is never forwarded to the PIM-SM domain.

The IGMP proxy is a way to propagate IGMP joins across router boundaries. The proxy triggers the boundary router connected to a PIM-SM domain to query for multicast flows and forward them to the PIM-DM domain. IGMP needs to be configured on all VLAN interfaces on which the proxy is to be forwarded or received and PIM-DM must be running for the traffic to be forwarded.

You can configure an IGMP proxy on a selected VLAN that will forward IP joins (reports) and IGMP leaves to the upstream border router between the two multicast domains. You must specify the VLANs on which the proxy is enabled as well as the address of the border router to which the joins are forwarded.

How IGMP Proxy Forwarding Works

The following steps illustrate how to flood a flow from the PIM-SM domain into the PIM-DM domain when an IGMP join for that flow occurs in the PIM-DM domain (refer to figure 2-4).

1. Routing Switch 1 is configured with the IGMP proxy forwarding function to forward joins towards Border Router 1. Routing Switch 1 is also configured to forward joins from VLAN 1 toward Border Router 2, as is VLAN 4 on Routing Switch 3.
2. VLAN 2 on Routing Switch 2 is configured to forward joins toward Border Router 1.
3. When the host connected in VLAN 1 issues an IGMP join for multicast address 235.1.1.1, the join is proxied by Routing Switch 1 onto VLAN 2 and onto VLAN 4. The routing information table in Routing Switch 1 indicates that the packet to Border Router 1 and Border Router 2 is on VLAN 2 and VLAN 4, respectively.

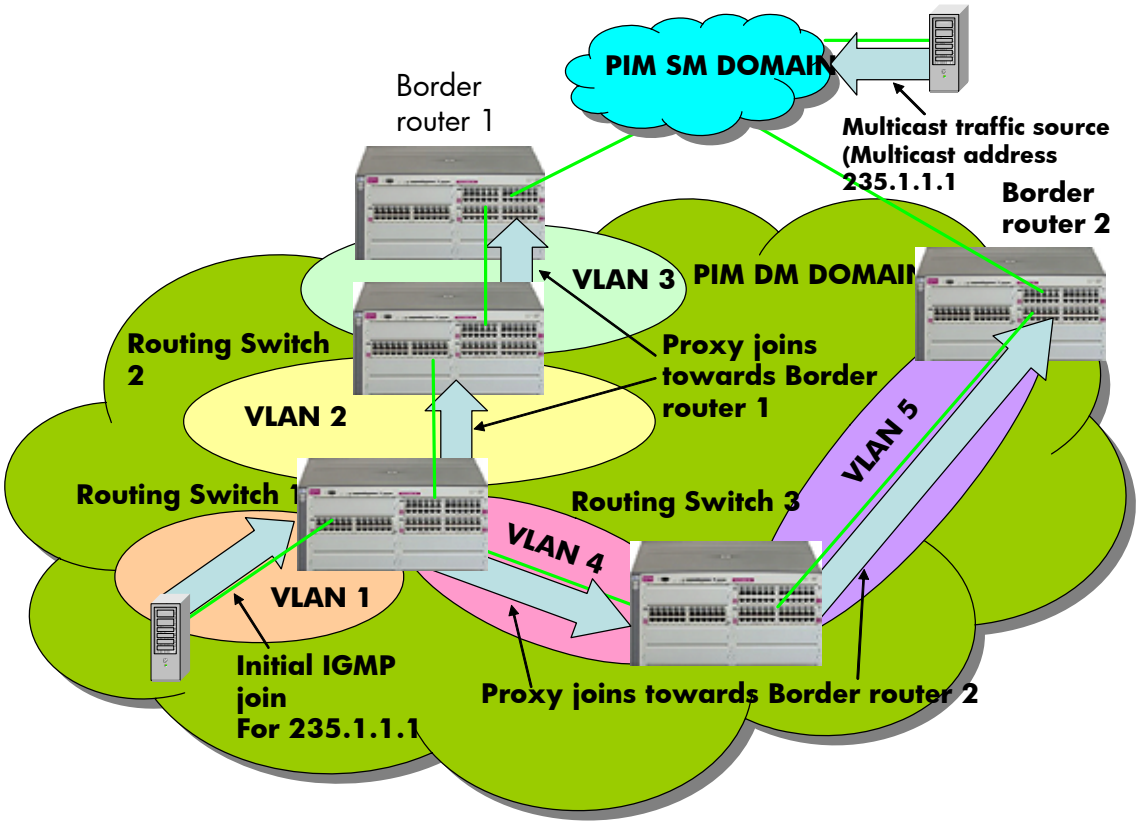

Figure 2-4. IGMP Proxy Example

4. Routing Switch 2 then proxies the IGMP join into VLAN 3, which is connected to Border Router 1.
5. Border Router 1 uses PIM-SM to find and connect to the multicast traffic for the requested traffic. The traffic is flooded into the PIM-DM network where it is routed to the original joining host.
6. Additionally, the join was proxied from Routing Switch 3 to Border Router 2. At first, both border routers will flood the traffic into the PIM-DM domain. However, PIM-DM only forwards multicasts based on the shortest reverse path back to the source of the traffic as determined by the unicast routing tables (routing FIB). Only one multicast stream is sent to the joining host. This configuration provides a redundant link in case the first link fails.

CLI Commands for IGMP Proxy Configuration

Syntax: [no] igmp-proxy-domain <domain-name> [<border-router-ip-address> <mcast-range | all>]

Add or leave a multicast domain. The **no** *form of the command is used to remove a multicast domain. All VLANs associated with the domain must first be removed for this command to work. See the* **no** *form of* **igmp-proxy** in the VLAN context command.

domain-name

User-defined name to associate with the PIM border router and multicast range that is being sent to toward the border router.

border-router-ip-addr

The IP address of the border router toward which IGMP proxy packets are sent. Not required for the **no** *form of the command.*

Note: The current routing FIB determines the best path towards the border router and therefore the VLAN that a proxy is sent out on.

<low-bound-ip-address | all>

The low boundary (inclusive) of the multicast address range to associate with this domain (for example, 234.0.0.1).

If **all** *is selected, the multicast addresses in the range of 224.0.1.0 - 239.255.255.255 will be included in this domain.*

Note: Addresses 224.0.0.0 - 224.0.0.255 are never used since these addresses are reserved for protocols.

<high-bound-ip-address>

The high boundary (inclusive) of the multicast address range to associate with this domain (for example, 236.1.1.1)

The following example shows the IGMP proxy border IP address (111.11.111.111) being configured.

```
ProCurve(config)# igmp-proxy-domain Bob 111.11.111.111
```

Figure 2-5. An example of the IGMP Proxy Border IP Address Command

The example below shows the lower and upper boundaries of the multicast address range associated with the domain named Bob.

```
ProCurve(config)# igmp-proxy-domain Bob 111.11.111.111 234.0.0.1
ProCurve(config)# igmp-proxy-domain Bob 111.11.111.111 236.1.1.1
```

Figure 2-6. Setting the Lower and Upper Bounds for Multicasting

VLAN Context Command

The following command is performed when in VLAN context mode. When a query occurs on the upstream interface, an IGMP join will be sent for all multicast addresses that are currently joined on the downstream interface.

Syntax: [no] igmp-proxy <domain-name>

Tells the VLAN which IGMP proxy domains to use with joins on the VLAN. The **no** *version of the command with no domain name specified removes all domains associated with this VLAN.*

Note: Multiple different domains may be configured in the same VLAN context where the VLAN is considered the downstream interface. The domain name must exist prior to using this command to add the domain.

Note

If the unicast routing path to the specified IP address was through the VLAN specified, then no proxy IGMP would occur, that is, a proxy is not sent back out on the VLAN that the IGMP join came in on.

If no unicast route exists to the border router, then no proxy IGMP packets will be sent.

IGMP Proxy Show Command

Syntax: show igmp-proxy < entries | domains | vlans >

Shows the currently active IGMP proxy entries, domains, or vlans.

```
ProCurve(config)# show igmp-proxy entries

  Total number of multicast routes: 2

  Multicast Address Border Address   VID  Multicast Domain
  ----------------- --------------  ----- ------
  234.43.209.12     192.168.1.1      1    George
  235.22.22.12      15.43.209.1      1    SAM
  226.44.3.3        192.168.1.1      2    George
```

Figure 2-7. Example Showing Active IGMP Proxy Entries

```
ProCurve(config)# show igmp-proxy domains


  Total number of multicast domains: 5

Multicast Domain  Multicast Range          Border Address   Active entries
 --------------- -------------------      ----------------  -----
 George           225.1.1.1/234.43.209.12   192.168.1.1       2
 SAM              235.0.0.0/239.1.1.1       15.43.209.1       1
 Jane             236.234.1.1/236.235.1.1   192.160.1.2       0
 Bill             ALL                       15.43.209.1       0
```

Figure 2-8. Example Showing IGMP Proxy Domains

```
ProCurve(config)# show igmp-proxy vlans

 IGMP PROXY VLANs

 VID      Multicast Domain   Active entries
 ------   ----------------   --------------
 1             George             1
 1             Sam                1
 1             Jane               0
 2             George             1
 4             George             0
 4             Bill               0
```

Figure 2-9. Example Showing Active IGMP Proxy VLANs

Operating Notes for IGMP Proxy Forwarding

- You can configure up to 12 multicast domains. These domains will indicate a range of multicast addresses and the IP address of the PIM-SM/PIM-DM border router.
- You must give each domain a unique name, up to 20 characters long.
- The domains may have overlapping multicast ranges.
- The IP address of the border router may be the same or different in each configured domain.
- Duplicate IGMP joins are automatically prevented, or leaves that would remove a flow currently joined by multiple hosts.
- Range overlap allows for redundant connectivity and the ability for multicasts to arrive from different border routers based on the shortest path back to the source of the traffic.
- The configured domain names must be associated with one or more VLANs for which the proxy joins are to be done.
- All routers in the path between the edge router receiving the initial IGMP packets and the border router have to be configured to forward IGMP using IGMP proxy.
- All upstream and downstream interfaces using IGMP proxy forwarding require IGMP and PIM to be enabled.
- You must remove all VLAN associations with the domain name before that domain name can be removed.

- The appropriate border routers must be used for each VLAN, or PIM-DM will not forward the traffic. This could occur when multiple border routers exist. It may be necessary to configure multiple overlapping domains if the multicast source address can generate the same multicast address and have different best paths to the PIM-DM domain.

Caution

Be careful to avoid configuring a IGMP forward loop, as this would leave the VLANs in a joined state forever once an initial join is sent from a host. For example, a join is issued from the host in VLAN 2 and routing switch 2 will proxy the join onto VLAN 1. Routing switch 3 will then proxy the join back onto VLAN 2 and increment its internal count of the number of joins on VLAN 2. Even after the host on VLAN 2 issues a leave, the proxy join will continue to remain and refresh itself each time a query occurs on VLAN 2. This type of loop could be created with multiple routers if an IGMP proxy is allowed to get back to the VLAN of the router that initially received the IGMP join from a host. (See figure 2-10.)

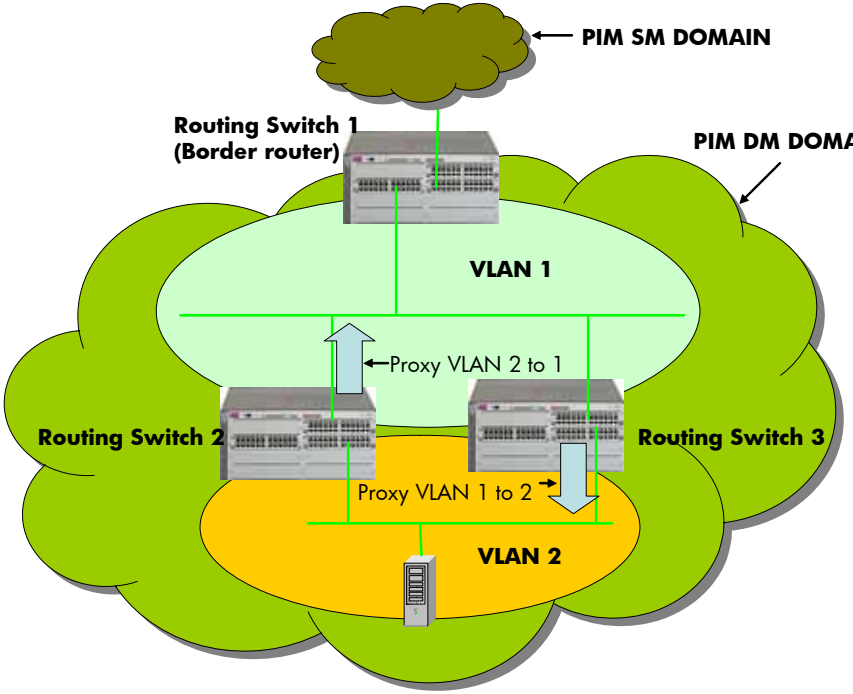

Figure 2-10. Proxy Loop Scenario

Using the Switch as Querier

The function of the IGMP Querier is to poll other IGMP-enabled devices in an IGMP-enabled VLAN to elicit group membership information. The switch performs this function if there is no other device in the VLAN, such as a multicast router, to act as Querier. Although the switch automatically ceases Querier operation in an IGMP-enabled VLAN if it detects another Querier on the VLAN, you can also use the switch's CLI to disable the Querier capability for that VLAN.

Note

A Querier is required for proper IGMP operation. For this reason, if you disable the Querier function on a switch, ensure that there is an IGMP Querier (and, preferably, a backup Querier) available on the same VLAN.

If the switch becomes the Querier for a particular VLAN (for example, the DEFAULT_VLAN), then subsequently detects queries transmitted from another device on the same VLAN, the switch ceases to operate as the Querier for that VLAN. If this occurs, the switch Event Log lists a pair of messages similar to these:

```
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: Other Querier detected
I 01/15/01 09:01:13 igmp: DEFAULT_VLAN: This switch is no longer Querie
```

In the above scenario, if the other device ceases to operate as a Querier on the default VLAN, then the switch detects this change and can become the Querier as long as it is not pre-empted by some other IGMP Querier on the VLAN. In this case, the switch Event Log lists messages similar to the following to indicate that the switch has become the Querier on the VLAN:

```
I 01/15/01 09:21:55 igmp: DEFAULT_VLAN: Querier Election in process
I 01/15/01 09:22:00 igmp: DEFAULT_VLAN: This switch has been elected
```

Excluding Well-Known or Reserved Multicast Addresses from IP Multicast Filtering

Each multicast host group is identified by a single IP address in the range of 224.0.0.0 through 239.255.255.255. Specific groups of consecutive addresses in this range are termed "well-known" addresses and are reserved for pre-defined host groups. IGMP does not filter these addresses, so any packets the switch receives for such addresses are flooded out all ports assigned to the VLAN on which they were received (except the port on which the packets entered the VLAN).

The following table lists the 32 well-known address groups (8192 total addresses) that IGMP does not filter on.

Table 2-2.IP Multicast Address Groups Excluded from IGMP Filtering

Groups of Consecutive Addresses in the Range of *224.0.0.X* to 239.0.0.*X**		**Groups of Consecutive Addresses in the Range of *224.128.0.X* to 239.128.0.*X****	
224.0.0.*x*	232.0.0.*x*	224.128.0.*x*	232.128.0.*x*
225.0.0.*x*	233.0.0.*x*	225.128.0.*x*	233.128.0.*x*
226.0.0.*x*	234.0.0.*x*	226.128.0.*x*	234.128.0.*x*
227.0.0.*x*	235.0.0.*x*	227.128.0.*x*	235.128.0.*x*
228.0.0.*x*	236.0.0.*x*	228.128.0.*x*	236.128.0.*x*
229.0.0.*x*	237.0.0.*x*	229.128.0.*x*	237.128.0.*x*
230.0.0.*x*	238.0.0.*x*	230.128.0.*x*	238.128.0.*x*
231.0.0.*x*	239.0.0.*x*	231.128.0.*x*	239.128.0.*x*

* *X* is any value from 0 to 255.

Notes:

IP Multicast Filters. *This operation applies to the ProCurve Series 5400zl switches, the Series 3500yl switches, the switch 6200yl, the switch 8212zl, the Series 5300xl switches, as well as the 1600M, 2400M, 2424M, 4000M, and 8000M, but not to the Series 2500, 2650, Series 4100gl, Series 4200vl, or 6108 switches (which do not have static traffic/security filters).*

IP multicast addresses occur in the range from 224.0.0.0 through 239.255.255.255 (which corresponds to the Ethernet multicast address range of 01005e-000000 through 01005e-7fffff). Where a switch has a static Traffic/Security filter configured with a "Multicast" filter type and a "Multicast Address" in this range, the switch will use the static filter unless IGMP learns of a multicast group destination in this range. In this case, IGMP dynamically takes over the filtering function for the multicast destination address(es) for as long as the IGMP group is active. If the IGMP group subsequently deactivates, the switch returns filtering control to the static filter.

Reserved Addresses Excluded from IP Multicast (IGMP) Filtering. Traffic to IP multicast groups in the IP address range of 224.0.0.0 to 224.0.0.255 will always be flooded because addresses in this range are "well known" or "reserved" addresses. Thus, if IP Multicast is enabled and there is an IP multicast group within the reserved address range, traffic to that group will be flooded instead of filtered by the switch.