Technical Guide

IGMP and MLD

Feature Overview and Configuration Guide

Introduction

Allied Telesis routers and managed Layer 3 switches use IGMP (Internet Group Management Protocol) and MLD (Multicast Listener Discovery) to track which multicast groups their clients belong to. This enables them to send the correct multimedia streams to the correct destinations. IGMP is used for IPv4 multicasting, and MLD is used for IPv6 multicasting.

This guide describes basic and advanced IGMP and MLD configuration, in the following major sections:

- an overview of IGMP/MLD and definitions of some of the IGMP/MLD terminology
- examples and discussion of the most common IGMP functionality—IGMP/MLD snooping, IGMP/MLD Querier behaviour and selection, and IGMP/MLD proxy
- examples and discussion of the advanced functionality available through the feature-rich IGMP/MLD implementation in AlliedWare Plus™
- information for debugging
- information about the STP state of the simple three-switch ring used in most examples
- information about IGMPv3/MLDv2

Products and software version that apply to this guide

This guide applies to all AlliedWare Plus™ products, running version 5.4.4 or later.

Feature support may change in later software versions. For the latest information, see the following documents:

- The product's Datasheet
- The product's Command Reference

Related documents

The following documents give more information about the multicasting features on AlliedWare Plus products:

- Protocol Independent Multicast - Sparse Mode (PIM-SM) Feature Overview and Configuration Guide
- PIM Sparse Mode for IPv6 (PIM-SMv6) Feature Overview and Configuration Guide
- PIM Dense Mode (PIM-DM) Feature Overview and Configuration Guide
- Command Reference for each product

These documents are available from the links above or on our website at alliedtelesis.com.

Content

Similarity of IGMP and MLD

MLD is the IPv6 equivalent of IGMP. Fortunately, as with most aspects of IPv6 routing, the features and operation of MLD closely resembles those of its IPv4 equivalent. In fact, the RFCs defining MLD quite explicitly state that they are adaptations of IGMP to IPv6. The packet types, timers, actions etc. in MLD are very much the same as those in IGMP. As with IGMP, there are multiple versions of MLD. Because MLD has been defined more recently than IGMP, the version numbers of MLD are actually one behind the version numbers of IGMP. This means that:

- MLDv1 (RFC2710) is equivalent to IGMPv2.
- MLDv2 (RFC3810) is equivalent to IGMPv3.

Because the protocols operate in such a similar manner, we describe them together in this single Feature Overview and Configuration Guide.

Terminology

Table 1: IGMP and MLD terms

TERM	DESCRIPTION
Multicast or Multicast stream	A flow of information, usually video or audio, that can go from one source to many destination clients.
Group	A multicast stream that clients can join. Groups have IPv4 addresses in the 224.0.0.0/4 range or IPv6 addresses in the FF00::/8 range.
Group member	A client that belongs to a particular multicast group. Which, in practical terms, means that it has sent out messages requesting it receive any stream that is destined to that group, and that, upon receiving such a stream, will process it in some way.
IGMP/MLD Querier or Designated Router	A device in a sub-network that is the coordinator for all multicast streams and IGMP/MLD membership information. Each sub-network has only one Querier. The Querier generates Membership Query messages to check which clients are group members, and processes Membership Reports and Leave messages.
IGMP/MLD Snooper	A device that spies on IGMP/MLD messages to create flow efficiencies by ensuring that multicast data streams are only sent to interested ports. A Snooper can decide on the best path to send multicast packets at Layer 2 but it cannot alter those packets or generate its own IGMP/MLD messages.
IGMP/MLD Proxy	A device that passes membership reports upstream towards a source in another subnet and multicast streams and queries downstream towards one or more downstream IP subnets. The proxy acts on behalf of clients and servers by altering packets.

IGMP/MLD Overview

Clients in an IP subnetwork use IGMP/MLD to indicate that they are interested in receiving a multicast stream. IGMP /MLD then ensures that routers and switches forward multicast packets out the appropriate ports to the interested clients.

The basic actions that occur in the IGMP/MLD protocol are:

1. Hosts that wish to receive certain multicast groups send out IGMP/MLD reports requesting the groups.
2. Routers and switches that receive these reports create forwarding entries to forward the groups in question out via the interface on which the report was received.
3. Routers send out IGMP/MLD queries at intervals, to find out who still wants to receive each group. The hosts who still want to receive the Group specified in any given query will respond with a new report, indicating their desire to continue receiving the group.
4. If a router or switch has not recently seen any reports requesting a given group on a given interface, it will remove the forwarding entry that is delivering that group to that interface.

In the actions above, a router is an IGMP/MLD querier. The other switches in the network (i.e. those between the listening hosts and the querier) will perform IGMP/MLD snooping.

The main duties of an IGMP/MLD querier are to:

- Send general queries, to find out who is still listening.
- Handle incoming reports, creating forwarding entries for the groups requested in those reports. If the IGMP/MLD querier is also acting as a Layer 3 router of IP multicast (using PIM or similar), the IGMP/MLD query process also needs to tell the Layer3 multicast routing process to send upstream requests for the groups being requested in incoming IGMP/MLD reports.
- Handle incoming Leave/Done messages. Removing the forwarding entry for the groups specified in the Leave/Done message, after first checking that no other host on the same interface is still listening to that group.

Strictly speaking, IGMP/MLD messages stay within a single IP subnet (although, a Proxy will transport IGMP/MLD messages from one subnet to another). So, the IGMP/MLD signaling is designed to control the flow of multicast within a subnet. The control of the flow of multicast between IP subnets is the job of a layer-3 multicast routing protocol like PIM.

Queriers and Snoopers

It is neither necessary nor desirable for every router or switch in an IP subnetwork to coordinate multicast traffic flows. Instead, a single router or switch does this and is called the **Querier** or the **Designated Router**. The Querier generates Query messages to check group membership, and processes Membership Reports and Leave messages.

However, other routers and switches in the network need to know whether to send multicasts out each of their ports. They find out this information by becoming **Snoopers**. Each Snooper checks

IGMP/MLD messages before forwarding them to and from the Querier, and uses the information in the messages to determine which ports to send multicasts out of.

The key differences between a network's Querier and its Snoopers are:

- The Querier generates Query messages to find out which ports need to transmit each multicast stream. The Snoopers also use Query messages to find this out, but they do this by passing on the Querier's messages—Snoopers cannot create Query messages themselves.
- The Querier has IGMP/MLD enabled as part of its IP configuration. Snoopers do not require any configuration because snooping is enabled by default on Allied Telesis routers and managed Layer 3 switches.
- Querying is a Layer 3 feature—the Querier looks into the IP headers of packets to determine whether to forward them. IGMP/MLD snooping is a Layer 2 feature. It does not require an IP configuration.

Messages

The following table describes the different IGMP/MLD messages in more detail:

Table 2: IGMPv2/MLDv1 message types

MESSAGE TYPE	DESCRIPTION
Membership Report	A client sends this when it wants to receive a multicast group. The Membership Report (sometimes called a join) is essentially a message that declares an interest in listening to a specified group.
Leave/Listener Done	A client sends this when it wants to leave a group. In IGMP, this is called a Leave message. In MLD, it is called a Listener Done message.
General Query	The Querier sends this to all clients—whether or not the Querier is currently sending multicasts to the client—to find out which groups they are listening to. Responses to General Queries ensure that the Querier's group membership information stays up to date. The group address field for IGMP General Queries is set to 0.0.0.0. They are sent to a destination address of 224.0.0.1, and by default Allied Telesis routers send them every 125 seconds. In the case of MLD, the Group address is ::, and the destination address is FF02::1.
Specific Query	The Querier sends this to a group address, to check whether clients are still listening to that group. The Querier sends a Specific Query after a client sends a Leave/Done message for that group. Specific Queries enable the Querier to confirm when all downstream clients have left a group, so that the Querier can stop sending the multicast stream.
Membership Query	This is a general term for both Specific and General Queries.
Query Solicit	Switches send this when STP or EPSR detects a topology change. The Querier responds by sending a General Query immediately instead of waiting until groups time out. This remaps IGMP/MLD to the new topology as quickly as possible.

As with a lot of the signaling packets in IPv6, MLD packets are ICMPv6 packets. The different types of MLD packets have different ICMPv6 type values, as shown in the table below.

Table 3: MLD packet types and ICMPv6 types

PACKET TYPE	SOURCE ADDRESS	DESTINATION ADDRESS	ICMP TYPE	MULTICAST ADDRESS IN PACKET
General query	Sender's link-local address	FF02: : 1	130	: :
Multicast address-specific query	Sender's link-local address	The specific group address being queried about	130	The specific group address being queried about
Report	Sender's link-local address	The address of the group the host wishes to receive	131	The address of the group the host wishes to receive
Listener Done	Sender's link-local address	FF02: : 2	132	The address of the group the host no longer wishes to receive

Choosing group addresses

This section describes things you need to be aware of when choosing addresses for your multicast groups.

Reserved IP addresses

IPv4 addresses in the range 224.0.0.0-239.255.255.255 and IPv6 addresses in the range FF00::/8 are multicast addresses, but many addresses in this range are **reserved**. Therefore, before choosing a multicast address, you should check its status in the "Internet Multicast Addresses" document at the IANA website at www.iana.org/assignments/multicast-addresses.

IPs using the same MAC

Another complication is that multicasting is designed to use each packet's group IP address to determine a multicast MAC address to send the packet to. However, multicasting does not have a 1:1 mapping of IP address to MAC address—instead each multicast MAC address corresponds to 32 multicast IPv4 addresses and for IPv6 there are 2^88 IPv6 groups per multicast MAC address. This means that different multicast IP addresses use the same MAC address.

- For IPv4, the multicast MAC address consist of 01-00-5E followed by the last 23 bits of the IP address.
- For IPv6, The MAC address for a multicast packet is created by adding the last 4 bytes of the IPv6 address to 3333.0000.0000. For example, if the IPv6 group address is: FF02:0000:0000:0000:0001: FF28:9C5A, then the MAC address for the group is: 3333.FF28.9C5A

For accurate control of multicast groups within a subnet, it is desirable to avoid using multiple IP addresses that have the same MAC address. In practice, this means that if you use x.0.y.z, then do not use x.128.y.z (or vice versa), where x is anything from 224-239, and y and z are the same in each IP address. For example, if y=6 and z=200 then these IP addresses use the same MAC: 224.0.6.200, 224.128.6.200, 225.0.6.200, 225.128.6.200, and so on.

To see this in detail, consider 224.0.6.200.

This has a multicast MAC of 01-00-5E-00-06-C8, like this:

IP address, decimal:	224.	0.	6.	200
IP address, binary:	11100000	00000000	00000110	11001000
MAC address, binary:		0000000	00000110	11001000
MAC address, hex:	01-00-5E	-00	-06	-C8

Therefore, the following multicast IP addresses will all have the same MAC address as 224.0.6.200, because their last 23 bits are all the same:

IP address, decimal:	IP address, binary:	
224.0.6.200	11100000 0	0000000 00000110 11001000
224.128.6.200	11100000 1	0000000 00000110 11001000
225.0.6.200	11100001 0	0000000 00000110 11001000
225.128.6.200	11100001 1	0000000 00000110 11001000
226.0.6.200	11100010 0	0000000 00000110 11001000
226.128.6.200	11100010 1	0000000 00000110 11001000
227.0.6.200	11100011 0	0000000 00000110 11001000
227.128.6.200	11100011 1	0000000 00000110 11001000
...	...	...
239.0.6.200	11101111 0	0000000 00000110 11001000
239.128.6.200	11101111 1	0000000 00000110 11001000
	Different IPs	The same MAC

Avoid x.0.0.y, x.0.1.y, x.128.0.y, and x.128.1.y

It is particularly important to avoid using any address in the ranges x.0.0.y, x.128.0.y, x.0.1.y, or x.128.1.y (where x is 224-239 and y is 1-254).

This is because x.0.0.y and x.128.0.y will map to the same multicast MAC address as 224.0.0.y. Similarly, x.0.1.y and x.128.1.y will map to the same multicast MAC address as 224.0.1.y. Most addresses in the ranges 224.0.0.y and 224.0.1.y are reserved for contacting all routers, or for routing protocol messages, so they are always flooded out all ports in the relevant VLAN.

Therefore, all addresses in the ranges x.0.0.y, x.128.0.y, x.0.1.y, or x.128.1.y are flooded out every port in the relevant VLAN. Using these addresses can significantly increase multicast traffic in your network.

If you are debugging a situation where it seems that certain multicast groups are forwarded when you think they shouldn't be, check whether the choice of group addresses has violated any of the recommendations above.

Limitation of MLD on AlliedWare Plus Switches

There is a 100 MLD interface limit when applying MLD commands to multiple VLANs. Only the first 100 VLANs have the required multicast structures added to the interfaces that allow multicast routing. MLD requires memory for storing data structures, as well as the hardware tables to implement hardware routing. As the number of ports, VLANs, static and dynamic groups increases then more memory is consumed. You can track the memory used for MLD with the command:

```
awplus# show memory pools nsm | grep
```

IGMP/MLD Snooping

In the example "Configuration example" on page 12, we discuss IGMP/MLD snooping, the key to efficient multicast traffic flow in a Layer 2 network.

Because IGMP and MLD snooping have minimal impact on system resources and are rarely themselves the source of networking issues, these features are enabled by default on switch ports in Allied Telesis managed Layer 3 switches and routers—they do not require any configuration.

It is possible to disable IGMP and MLD snooping, including on a per-VLAN basis, using the commands **no ip igmp snooping** and **no ipv6 mld snooping**. However, we recommend leaving both IGMP and MLD snooping enabled on switches unless an Allied Telesis support representative tells you to disable them, even if you are not actively using the switch for multicast.

In a single-switch network, IGMP/MLD snooping makes multicasting happen with no configuration at all. All you need to do is connect your server and clients to the switch.

In a multi-switch network, at least one switch must also have an IGMP/MLD configuration. This switch is called the IGMP/MLD Querier and coordinates the flow of multicast information through the network. The following example describes a multi-switch configuration, so as well as discussing the effect of IGMP/MLD snooping, it outlines the actions that the Querier takes. "Explanation of multiple potential IGMP/MLD Queriers" on page 23 discusses the role of the Querier in greater detail.

Explanation of IGMP/MLD Snooping

This section steps through the events that occur in a typical use of multicasting in this network, i.e. to stream multicast packets for a group. The following figure shows the process by which IGMP/MLD tracks multicast clients and ensures that the correct clients receive the stream.

What happens before a multicast client exists:

1. Querier starts receiving multicast stream from server. Querier has no interested clients so doesn't forward multicast.

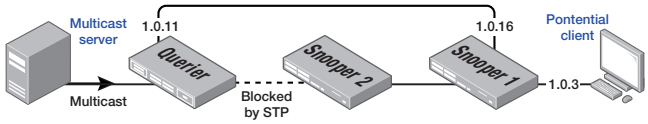

2. Querier sends General Query (GQ) to find out if any would-be clients exist. Snooper 1 receives Query on port 1.0.16, snoops Query, and creates All Groups entry for port 1.0.16.

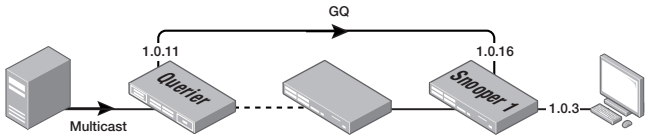

3. Snooper 1 forwards Query out all ports. No clients exist, so no clients reply to Query.

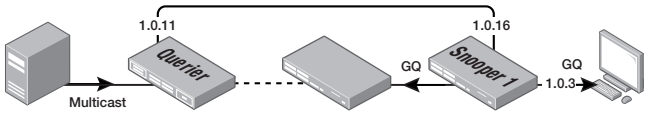

4. Client joins group by sending Membership Report to the group address.

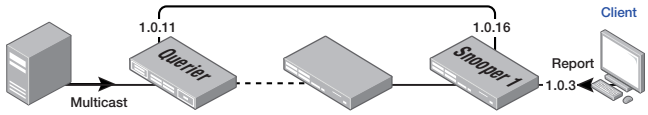

5. Snooper 1 receives Report on port 1.0.3, snoops Report, and adds a group entry for port 1.0.3. Snooper 1 forwards Report out its All Groups port.

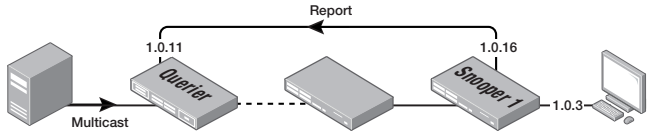

6. Querier receives Report on port 1.0.11 and adds a group entry for port 1.0.11. Querier starts forwarding multicast stream out port 1.0.11.

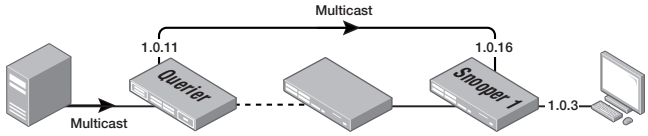

7. Snooper 1 receives multicast stream and forwards it out port 1.0.3.

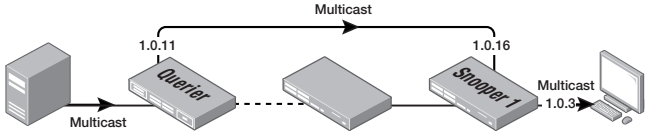

8. Querier continues to send General Queries periodically. These keep All Groups entries alive on Snoopers.

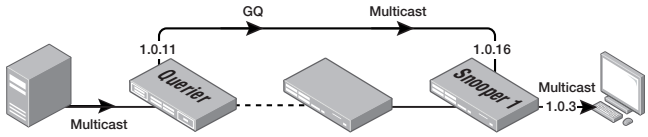

Configuration example

This example has a three-switch loop, as shown in the following diagram. Switch-1 (Core) is running IGMP/MLD and the other two switches (Edge) are running IGMP/MLD snooping.

In AlliedWare Plus (AW+), both IGMP/MLD snooping and RSTP are enabled by default. On the Core and Edge switches, the only configuration needed is to set the client-facing ports as spanning tree edge ports.

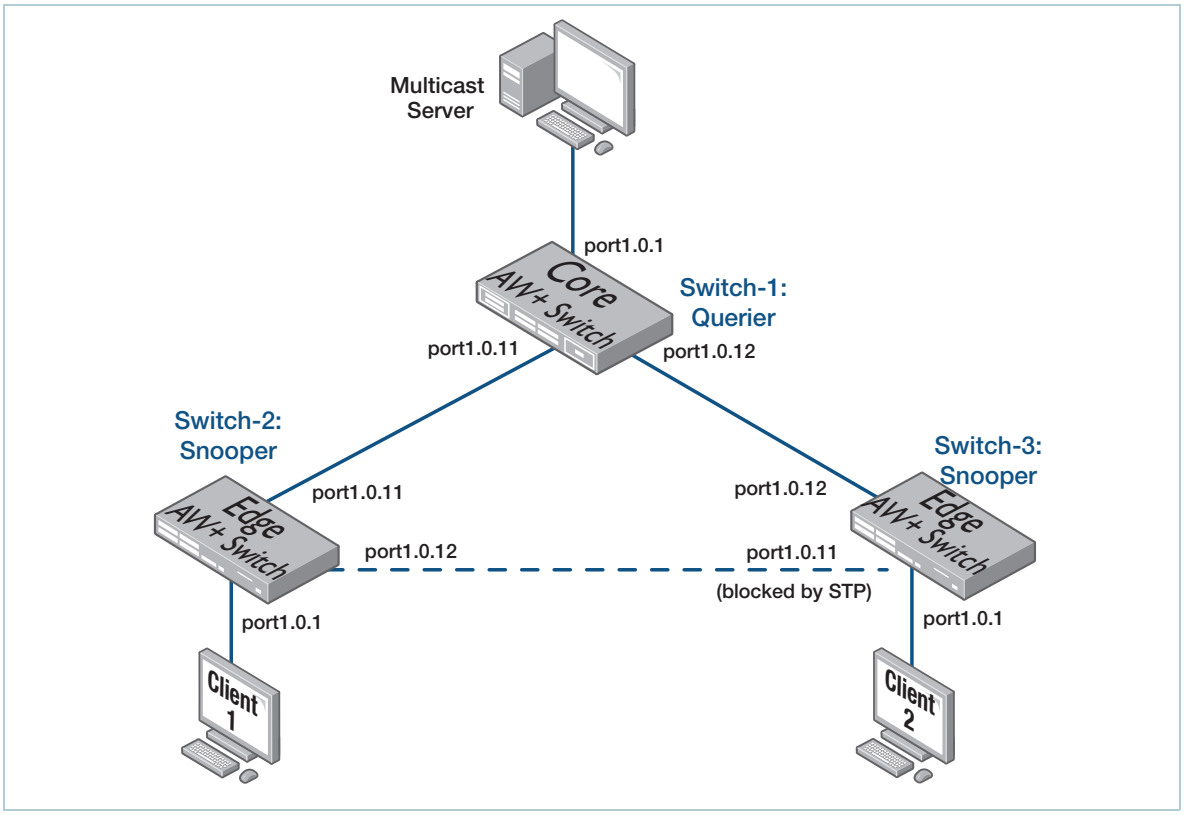

Step 1. **Configure Switch-1 (Core)**

Switch-1 is configured with IGMP and MLD which makes it the IGMP/MLD Querier in this network. It is best practice to make the Querier the closest switch to the multicast source, and in this example Switch-1 is closest. For more information about Queriers see "Multiple Potential IGMP/MLD Queriers" on page 20.

In this example, we add an IP interface to vlan100 and enable IP IGMP and IPv6 MLD on this interface so that Switch-1 becomes the IGMP/MLD Querier in this network.

The default version of IGMP in AlliedWare Plus is IGMPv3. In this example we change it to IGMPv2 to match the version the clients are using. Similarly, the default MLD version in AlliedWare Plus is MLDv2, so the example sets the MLD version to 1, to match the version being used by the clients.

```
!
hostname Switch-1
!
spanning-tree mode rstp
ipv6 forwarding
!
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
  switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
  switchport access vlan 100
 spanning-tree edgeport
!
interface vlan100
 ip address 192.168.100.254/24
 ip igmp
 ip igmp version 2
 ipv6 address 2001:0db8::12:252/64
 ipv6 enable
 ipv6 mld
 ipv6 mld version 1
```

AlliedWare Plus provides two ways of making this switch the Querier: as above, or by entering the following commands:

```
awplus# configure terminal
awplus(config)# int vlan100
awplus(config-if)# ip igmp snooping querier
awplus(config-if)# ipv6 mld snooping querier
```

These commands enable IGMP/MLD Querier operation on a VLAN when no multicast routing protocol is configured in the VLAN. When enabled, the IGMP/MLD Snooping Querier sends out periodic IGMP/MLD queries for all interfaces on that VLAN. This command applies to interfaces configured for IGMP/MLD Snooping.

This Snooping Querier function is a useful way to provide a Querier in a Layer 2 network that does not contain a multicast router. We do not recommend enabling the IGMP/MLD Snooping Querier feature on a Layer 2 switch when there is a router acting as an operational IGMP/MLD Querier in the network.

In terms of IGMP/MLD querying activity, there is no real difference between the Querier and Snooping Querier, they both send queries on the specified interface.

Step 2. **Configure Switch-2 (Edge)**

Switch-2 is an IGMP/MLD Snooper. It forwards multicast packets and IGMP/MLD messages as required. IGMP/MLD snooping is enabled by default and does not need any configuration.

```
!
hostname Switch-2
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
  switchport access vlan 100
  spanning-tree edgeport
!
interface port1.0.11-1.0.12
  switchport access vlan 100
!
interface port1.0.13-1.0.24
  switchport access vlan 100
  spanning-tree edgeport
```

Step 3. **Configure Switch-3 (Edge)**

Switch-3 is also an IGMP/MLD Snooper. It forwards multicast packets and IGMP/MLD messages as required. IGMP/MLD snooping is enabled by default and does not need any configuration.

```
!
hostname Switch-3
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
  switchport access vlan 100
  spanning-tree edgeport
!
interface port1.0.11-1.0.12
  switchport access vlan 100
!
interface port1.0.13-1.0.24
  switchport access vlan 100
  spanning-tree edgeport
```

Using the show command output to investigate IGMP state

No group members

At first we have no hosts requesting any multicast. The multicast server is streaming group 224.12.13.14 to the Querier, Switch-1. Switch-1 knows about the group, but has nobody interested in receiving it. You can see this by using the command **show ip igmp snooping statistics interface vlan100** on Switch-1. The output of this command shows that Switch-1 has an entry for the group, but no associated ports.

```
Switch-1#sh ip igmp snooping statistics int vlan100

IGMP Snooping statistics for vlan100
Group Type:     UnReg MC Group
Interface:      vlan100
Group:          224.12.13.14
Uptime:         00:00:10
Group mode:     Include (Expires: 00:04:10)
Last reporter:  192.168.100.100
```

No ports are mentioned in the output. Also, you can see that the Group Type is described as ‘UnReg MC Group’. This means that it is an entry for an unregistered group. That is a group for which data packets are arriving, but no IGMP reports have been received.

Client joins the group

When a client joins the group, the Group List changes for the Snooper that the client is attached to, and on the Querier. First, look at the output of the command **show igmp snooping statistics interface vlan100** on the Snooper.

```
Switch-2#sh ip igmp snooping statistics interface vlan100
IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         00:29:24
Group mode:     Include ()
Last reporter:  192.168.100.254

Port member list:
port1.0.11 - 169 secs

Interface:      vlan100
Group:          224.12.13.14
Flags:
Uptime:         00:00:30
Group mode:     Exclude (Expires: 00:03:50)
Last reporter:  192.168.100.10
Source list is empty

Port member list:
port1.0.1 - 230 secs
```

This output now shows two entries, one for each of the following:

- Group 224.12.13.14 and port 1.0.1. This shows that the client is attached to the Snooper through port 1.0.1 and is listening to group 224.12.13.14. The Snooper created this entry at stage 5 in the

process "Explanation of IGMP/MLD Snooping" on page 11. This entry means that the Snooper forwards packets from 224.12.13.14 out port 1.0.1.

- Group Type: Router port learnt on port 1.0.11. This shows that the Snooper is connected to the Querier through port 1.0.11. The Snooper created this entry at stage 2 in the process. This entry means that the Snooper forwards IGMP reports and Leave messages out port 1.0.11. More information about Router Ports is provided in the section "Router ports" on page 18.

Let us take the opportunity here to examine the information that appears in the output of the **show ip igmp snooping statistics interface** command.

- **Group Type**—This item only appears in entries that are Router Port or Unregistered group entries. For standard entries created for IGMP reports, this item does not appear.
- **Interface**—The interface which the IGMP reports were received on. In the case of a Router Port entry, the VLAN towards the router. For unregistered groups it is the VLAN the stream is arriving on.
- **Group**—The multicast group address.
- **Uptime**—How long the entry has been in existence
- **Group mode**—This parameter can either be Include or Exclude. The terms Include and Exclude come from IGMPv3 terminology. IGMPv3 has the concept of including sources or excluding sources that the group can be received from. If a group entry is created by IGMPv2 reports, then the entry will be labelled as **Exclude**, because IGMPv2 reports are treated as being “Exclude no source” reports. A Router Port entry is labelled as **Include** as this entry effectively includes no sources.
- **Last reporter**—The IP address of the last member to send an IGMP report for this group. In the case of a router port, it is the address of the router or Querier that the router-indicating packets are coming from. In the case of an unregistered group, it is the source address of the stream.

There is an equivalent command for IPv6:

`show ipv6 mld snooping statistics interface <interface name>`

The Querier receives the Report on port 1.0.11. Next, look at the output of the command **show ip igmp groups** on the Querier.

```
Switch-1# sh ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Enabled, Active, Querier, Configured for version 2
 Internet address is 192.168.100.254
```

```
Switch-1# sh ip igmp groups

IGMP Connected Group Membership
Group Address    Interface        Uptime   Expires   Last Reporter
224.12.13.14     vlan100          00:08:17 00:03:31  0.0.0.0
```

The output above shows an entry for group 224.12.13.14. This entry shows that the Querier knows about a client for 224.12.13.14. The Querier created this entry at stage 6 in the process. The fact that the **Last Reporter** in this entry is 0.0.0.0 indicates that the report that created this entry was not received from a directly connected host (the report was received via Switch-2). Because the report was forwarded by a snooper that has no IP address, the report arrived with IP source address 0.0.0.0.

The equivalent commands for IPv6 are:

```
show ipv6 mld interface
show ipv6 mld groups
```

Finally, look at the output of the command **show ip igmp snooping statistics interface vlan100** on the Querier. Even though Switch-1 is the Querier for this network instead of a Snooper, this command shows that a client for group 224.12.13.14 is reached out port 1.0.11. So a Querier does also perform snooping, to work out the exact egress port(s) for each group.

```
Switch-1# show ip igmp snooping statistics int vlan100

IGMP Snooping statistics for vlan100
Interface:      vlan100
Group:          224.12.13.14
Uptime:         00:12:11
Group mode:     Exclude (Expires: 00:03:43)
Last reporter:  0.0.0.0
Source list is empty

Port member list:
port1.0.11 - 224 secs
```

When a client leaves a group

When a client wants to stop receiving a group's multicast stream, it sends an IGMP Leave message or an MLD Listener Done message, with a destination address of the group. The Snooper forwards the Leave/Done message out its All Groups port, so the message arrives at the IGMP/MLD Querier. At this point, the IGMP/MLD Querier sends a series of Specific Queries (two by default) to see if anybody else is still listening to this group.

If the Snooper receives a response to the Specific Queries, it forwards the response to the Querier and continues to forward the multicast stream to the ports that want to receive it. If the Snooper does not receive a response to the Specific Queries, it stops forwarding the stream.

For a detailed description of how the leave process works, see "How Clients Leave Groups: Queries and Timers" on page 61.

Router ports

As mentioned a few times above, the port via which packets from a querier are arriving at a snooper is referred to as a Router Port. This is because this port is the port directed towards the local IGMP/MLD router.

To efficiently use bandwidth, multicasting needs to impose a tree structure onto a Layer 2 network, so there is a single unique path that runs from any given source to any given receiver. There needs to be a single switch that acts as the root of this tree. This root switch is the Querier.

All snoopers need to inform the Querier about the groups that they know hosts have joined, and must make sure that the Querier knows about any multicast streams that are being transmitted into the network.

Because of this, all switches will learn which port connects them back to the Querier. This port is referred to as the Router Port or All Groups Port.

The rules are:

- IGMP reports and leaves received from downstream are retransmitted out through the Router Port.
- The Router Port is added to the list of egress ports for ALL the groups that the switch is forwarding, so that all streams that arrive from anywhere (except those arriving on the Router Port) are forwarded toward the Querier.

Preventing occasional brief flooding of unregistered streams

The forwarding of multicast streams by switches is controlled by the entries that IGMP puts into the hardware tables.

If the switch has received reports that request the forwarding of a stream, then the switch will have hardware entries in place for sending that stream to the ports on which the reports have been received, whether the stream is arriving at the switch or not. Such streams are not flooded.

However, if a stream arrives at the switch and the switch has not received any IGMP reports from clients requesting that stream, the situation is a little different. Such streams are referred to as unregistered streams (see "No group members" on page 15). For such streams, the switch's default behavior is typically to flood the stream to all the ports in the VLAN on which the stream was received. This default behavior is overridden by IGMP snooping, which creates a hardware entry specifically to stop the unregistered stream from being flooded.

However, there are brief moments when this prevention is not in place, and an unregistered stream may be briefly flooded. This happens in the following two situations:

- When the stream first arrives at the switch
- When IGMP snooping removes the 'forward to no ports' entry, which it does periodically so that it can check if the stream is still arriving.

There is a command available to stop this flooding during even those brief periods. The command works by changing the default behavior of the switch chip from flooding the stream to sending the stream to the CPU instead. The command is:

```
awplus# platform stop-unreg-mc-flooding
```

This command is important in networks that contain equipment that is sensitive to receiving unsolicited multicast streams, for example due to the equipment having a low-powered CPU that is overwhelmed by the arrival of too much data, or due to some security-checking that reports problems if unexpected streams are received, or due to other reasons.

Please note the following:

- You should not use this command within any IPv6 networks, because it inhibits IPv6 neighbor discovery operation.
- This command does not affect the flooding of Local Network Control Block IPv4 multicast packets in the address range 224.0.0.1 to 224.0.0.255 (224.0.0/24). Such packets will continue to be uninterruptedly flooded, as they need to be.

Limitations of IGMP snooping on x210 and x230 Series switches

The x210 and x230 Series switches do not support SSM (Source Specific Multicast). When using IGMPv3 on these switches, make sure you use ASM (Any Source Multicast).

Flooding unregistered multicast packets to all ports without mirroring to CPU

Multicast packets from an unregistered source are normally mirrored to the CPU to ensure IGMP and PIM-SM knows about the group. In certain networks, it is possible for a large number of packets with a number of different sources destined for the same group address, to overwhelm a switches hardware table. This could cause packets to be stuck mirrored to the CPU forever, resulting in high CPU usage and in some cases stack failovers.

There is a command available that allows you to configure a static multicast group that will flood matching packets to all ports in the VLAN and not mirror any packets matching that group to the CPU. The command is:

```
awplus# ip igmp flood-group <ip-address> [<vlan-id>]
```

This command adds an all sources multicast entry into the switches multicast hardware table to flood multicast packets to all ports within the VLAN without mirroring the traffic to CPU. This significantly reduces the number of hardware entries consumed.

The Layer 3 variant of this command is only supported on one VLAN group address. Any number of the Layer 2 variants can be used.

Examples To configure an IGMP flooding group to L2 ports only, use the following commands. This will flood any UDP packet to group 239.255.255.250 to all ports in vlan1:

```
awplus# configure terminal
awplus(config)# int vlan1
awplus(config-if)# ip igmp flood-group 239.255.255.250
```

To configure an IGMP flooding group to L2 ports and L3 forwarding, use the following commands. This will flood any UDP packet to group 239.255.255.250 to all ports in vlan2 and forward any UDP packet to all ports in vlan3:

```
awplus# configure terminal
awplus(config)# int vlan2
awplus(config-if)# ip igmp flood-group 239.255.255.250 vlan3
```

To remove an IGMP flooding group from vlan2, use the following commands:

```
awplus# configure terminal
awplus(config)# int vlan2
awplus(config-if)# no ip igmp flood-group 239.255.255.250
```

IGMP flooding groups are integrated into the existing command: **show ip igmp groups**

Output Figure 1: Example output from **show ip igmp groups**

```
awplus#show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface   Uptime    Expires  Last Reporter
239.255.255.250  vlan3000    01:20:59  stopped  0.0.0.0
```

Multiple Potential IGMP/MLD Queriers

To find out more about IGMP/MLD we next investigate what happens when more than one router has an IGMP/MLD configuration. RFC 2236, Internet Group Management Protocol, Version 2, and RFC2710, Multicast Listener Discovery (MLD) for IPv6, states that each Layer 2 network should have only one IGMP/MLD Querier. You may configure IGMP/MLD on more than one router, perhaps for redundancy, but the routers have a pseudo-election and the device with the lowest IP becomes the operating IGMP/MLD Querier.

This pseudo-election; in reality, it is a query suppression, in that all queriers except that with the lowest IP address just stop sending out queries when they realize that the network contains a querier with a lower IP address than themselves. Each querier will initially send out regular queries. When a querier receives a query message with a lower IP address then it will stop sending its own queries for the duration of the Other Querier interval. As long as the other querier keeps regularly sending queries, the higher IP address querier will not send queries.

A snooping querier (with source address of 0.0.0.0/: :) should consider all real IP addresses to be ‘lower’ and stop sending queries when it receives queries from a non-snooping querier. The following example describes a network with two potential Queriers.

Example

The network for this example uses the same loop as for "IGMP/MLD Snooping" on page 10 and is shown in the following figure.

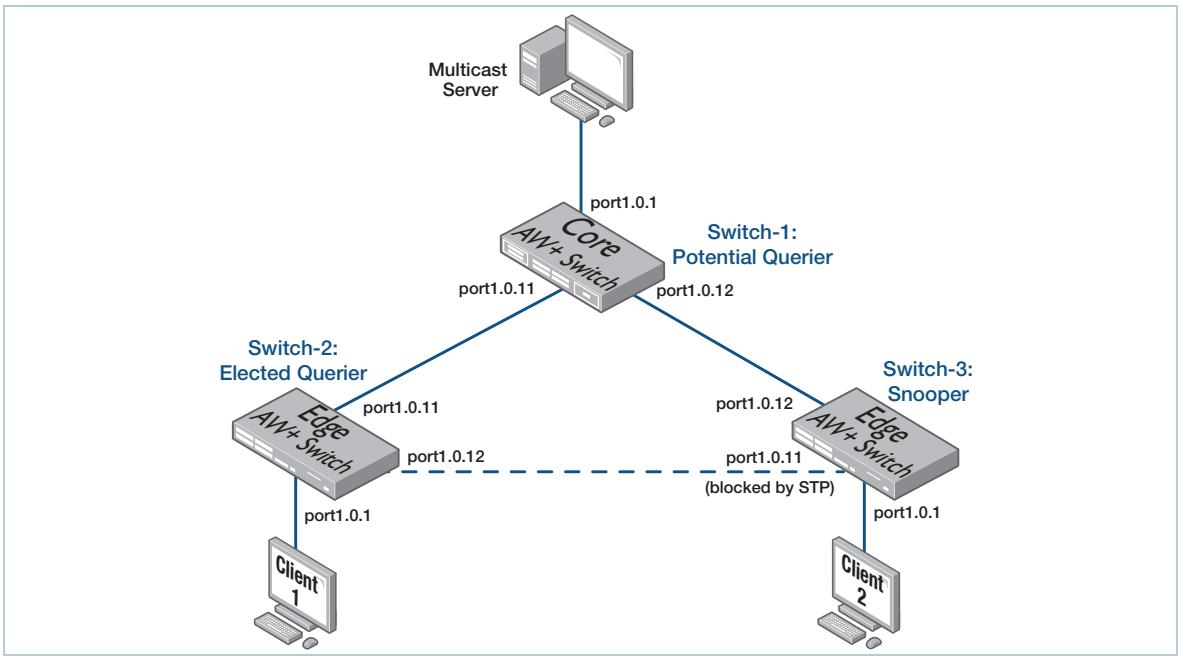

Both Switch-1 and Switch-2 are configured with IGMP/MLD, making both of them potential Queriers. Switch-3, by default configuration, is an IGMP/MLD Snooper.

Step 1. **Configure Switch-1**

Switch-1 is a potential IGMP/MLD Querier. It acts as a Snooper if not elected as the Querier.

```
!
hostname Switch-1
!
spanning-tree mode rstp
ipv6 forwarding
vlan database
 vlan 100 name vlan100
 !
interface port1.0.1-1.0.10
  switchport access vlan 100
  spanning-tree edgeport
!
interface port1.0.11-1.0.12
  switchport access vlan 100
!
interface port1.0.13-1.0.24
  switchport access vlan 100
  spanning-tree edgeport
!
interface vlan100
 ip address 192.168.100.254/24
 ip igmp
 ip igmp version 2
 ipv6 address 2001:0db8::12:252/64
 ipv6 enable
 ipv6 mld
 ipv6 mld version
```

Step 2. **Configure Switch-2**

Switch-2 is a potential IGMP/MLD Querier. It acts as a Snooper if not elected as the Querier.

```
!
hostname Switch-2
!
spanning-tree mode rstp
ipv6 forwarding
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
  switchport access vlan 100
  spanning-tree edgeport
!
interface port1.0.11-1.0.12
  switchport access vlan 100
!
interface port1.0.13-1.0.24
  switchport access vlan 100
  spanning-tree edgeport
!
interface vlan100
 ip address 192.168.100.253/24
 ip igmp
 ip igmp version 2
 ipv6 address 2001:0db8::12:251/64
 ipv6 enable
 ipv6 mld
 ipv6 mld version 1
```

Step 3. **Configure Switch-3**

Switch-3 is an IGMP/MLD Snooper. It forwards multicast packets and IGMP/MLD messages as required. IGMP/MLD snooping and Rapid Spanning Tree Protocol are enabled by default, so it does not need any specific IGMP/MLD or RSTP configuration. (The configuration is the same as the configuration used in the first example).

```
!
hostname Switch-3
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
  switchport access vlan 100
  spanning-tree edgeport
!
interface port1.0.11-1.0.12
  switchport access vlan 100
!
interface port1.0.13-1.0.24
  switchport access vlan 100
  spanning-tree edgeport
```

Explanation of multiple potential IGMP/MLD Queriers

When there are no group members

Switch-1 and Switch-2 are both possible Queriers, and an election determines which switch becomes the actual Querier. We can see the results of the election by using the command **show ip igmp** or **show IPv6 MLD interface** on each switch.

When there are two or more possible Queriers on the same network, the election process is based on the lowest IP address. So, in this example Switch-2, which has the IPv4 address of 192.168.100.253, will be the Designated Querier as it has the lowest IPv4 address. Similarly, Switch-2 has the lower IPv6 address, 2001:0db8::12:251/64, and so will be the designated MLD querier.

On Switch-1, we see that it reports itself as the non-querier, and shows that the Querier is 192.168.100.253, which is Switch-2.

```
Switch-1# sh ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Enabled, Active, Non-Querier, Configured for version 2
 Internet address is 192.168.100.254
 IGMP interface has 3 group-record states
 IGMP activity: 93 joins, 0 leaves
  IGMP querying router is 192.168.100.253
 IGMP robustness variable is 2
 IGMP last member query count is 2
 IGMP query interval is 125 seconds
 IGMP Startup query interval is 31 seconds
 IGMP Startup query count is 2
 IGMP query holdtime is 500 milliseconds
 IGMP querier timeout is 255 seconds
 IGMP max query response time is 10 seconds
 Last member query response interval is 1000 milliseconds
 Group Membership interval is 260 seconds
 IGMP Last member query count is 2
 Strict IGMPv3 ToS checking is disabled on this interface
 Source Address checking is enabled
 IGMP Snooping is globally enabled
 IGMP Snooping query solicitation is globally enabled for Root/Master Nodes
 IGMP Snooping query solicitation is globally disabled for Non Root/Master Nodes
  Num. query-solicit packets: 0 sent, 0 recvd
 IGMP Snooping is enabled on this interface
 IGMP Snooping fast-leave is not enabled
 IGMP Snooping querier is not enabled
 IGMP Snooping report suppression is enabled
```

Now we look at Switch-2, which has the lower IP address configured, resulting in it being the Designated Querier based on the IGMP Querier election process.

```
Switch-2#sh ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Enabled, Active, Querier, Configured for version 2
 Internet address is 192.168.100.253
 IGMP interface has 3 group-record states
 IGMP activity: 18 joins, 1 leaves
 IGMP robustness variable is 2
 IGMP last member query count is 2
 IGMP query interval is 125 seconds
 IGMP Startup query interval is 31 seconds
 IGMP Startup query count is 2
 IGMP query holdtime is 500 milliseconds
 IGMP querier timeout is 255 seconds
 IGMP max query response time is 10 seconds
 Group Membership interval is 260 seconds
 IGMP Last member query count is 2
 Last member query response interval is 1000 milliseconds
 Strict IGMPv3 ToS checking is disabled on this interface
 Source Address checking is enabled
 IGMP Snooping is globally enabled
 IGMP Snooping query solicitation is globally enabled for Root/Master Nodes
 IGMP Snooping query solicitation is globally disabled for Non Root/Master Nodes
  Num. query-solicit packets: 0 sent, 0 recvd
 IGMP Snooping is enabled on this interface
 IGMP Snooping fast-leave is not enabled
 IGMP Snooping querier is not enabled
 IGMP Snooping report suppression is enabled
```

The highlighted text shows that this switch has become the Designated IGMP Querier for vlan100. The Source now sends a stream onto the LAN. With no clients requesting the stream, we look at the IGMP Snooping statistics on Switch-1. We can see that the router port has been learnt from the IGMP Querier (Switch-2). There are no ports associated with the multicast group 224.12.13.14 as no users have requested this group. Again, this appears with Group Type of “UnReg MC Group”

```
Switch-1#sh ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         00:24:44
Group mode:     Include ()
Last reporter:  192.168.100.253

Port member list:
port1.0.11 - 168 secs

Group Type:     UnReg MC Group
Interface:      vlan100
Group:          224.12.13.14
Uptime:         00:00:13
Group mode:     Include (Expires: 00:04:07)
Last reporter:  192.168.100.100
```

From Switch-2, the IGMP Querier, we can see from the output that it is receiving this group, but there are no interfaces associated with this, as no clients have requested that stream. So it appears as an “UnReg MC Group”.

```
Switch-2#sh ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Interface:      vlan100
Group Type:     UnReg MC Group
Interface:      vlan100
Group:          224.12.13.14
Uptime:         00:03:58
Group mode:     Include (Expires: 00:00:21)
Last reporter:  192.168.100.100
```

The reason that the stream arrives at Switch-2 is because Switch-1 will forward **all** multicast out its Router Port. More information about Router Ports is provided in the section "Router ports" on page 18.

When a client joins a group

Now imagine that Client 2 sends a Membership Report to Switch-3 for the group 224.12.13.14. If we check the group membership for Switch-2 by using the command **show ip igmp snooping statistics interface vlan100**, we see a group entry for 224.12.13.14. It has changed from being an “UnReg MC” group entry to a normal forwarding entry.

We can see that this membership report from the client (192.168.100.20/24), which is attached to port1.0.1 on Switch-3, has been received on port1.0.11 on Switch-2.

In this case, we see that the **Last reporter** in the entry on Switch-2 is 192.168.100.254. This is because the report was forwarded to Switch-2 via Switch-1. Given that Switch-1 has the IP address 192.168.100.254 on vlan100, it put that IP address as the source address of the IGMP report when it forwarded it.

```
Switch-2#sh ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Interface:      vlan100
Group:          224.12.13.14
Flags:
Uptime:         00:01:49
Group mode:     Exclude (Expires: 00:03:08)
Last reporter:  192.168.100.254
Source list is empty

Port member list:
port1.0.11 - 189 secs
```

On Switch-3, we can see that reports have been received from a client (192.168.100.20) which is attached to port 1.0.1:

```
Switch-3#sh ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         00:13:02
Group mode:     Include ()
Last reporter:  192.168.100.253

Port member list:
port1.0.12 - 204 secs

Interface:      vlan100
Group:          224.12.13.14
Flags:
Uptime:         00:03:19
Group mode:     Exclude (Expires: 00:03:24)
Last reporter:  192.168.100.20
Source list is empty

Port member list:
port1.0.1 - 204 secs
```

If we check the group membership for Switch-1 and Switch-3, we see there are entries for 224.12.13.14, but also see a Router Port entry on each switch. The Router Port entry points to the Querier, Switch-2. The output for Switch-1, for example, shows port 1.0.11 as the Router port, indicating that Switch-1 reaches the Querier via port 1.0.11. More information about Router Ports is provided in the section "Router ports" on page 18.

Switch-1 has received the Membership report on port1.0.12 (this is the connection to Switch-3). We see port1.0.12 on the port member list of Switch-1 because interface port1.0.11 on Switch-3 is in a RSTP Discarding state (this is the connection to Switch-2.) The IGMP packets forwarded by Switch-3 are going to be received on interface port1.0.12 on Switch-1.

```
Switch-1#sh ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         00:41:56
Group mode:     Include ()
Last reporter:  192.168.100.253

Port member list:
port1.0.11 - 136 secs

Interface:      vlan100
Group:          224.12.13.14
Uptime:         00:06:51
Group mode:     Exclude (Expires: 00:02:18)
Last reporter:  0.0.0.0
Source list is empty

Port member list:
port1.0.12 - 139 secs
```

IGMP Proxy

In very simple tree-design networks, IGMP Proxy is a useful simple alternative to a multicast routing protocol for multicasting between VLANs.

An IGMP Proxy sends IGMP Membership Report and Leave group messages to an upstream sub-network on behalf of downstream devices, and sends Queries downstream. In other words, an IGMP Proxy effectively ferries IGMP messages from one VLAN to another. The IGMP Proxy looks like an IGMP Querier to the downstream VLAN, and like a client to the upstream VLAN. Note that the Proxy can only have one configured upstream VLAN, but it can service multiple downstream VLANs.

Note: AlliedWare Plus does not support MLD Proxy, only IGMP Proxy. IGMP proxy does not work with other multicast routing protocols, such as PIM-SM or PIM-DM.

The Proxy packets forward at line speed, and the switch creates hardware entries as it receives the IGMP requests and multicast data. This means that multicast data packets are not sent to the CPU for forwarding.

Example

IGMP Proxy only works in tree networks, so for this example we convert the network from a loop into a tree by disabling port 1.0.11 on Switch-1.

Switch-3 is the IGMP Proxy. Switch-1 is upstream of the Proxy. Switch-2 is downstream of the Proxy on vlan100. Therefore, the multicast server and Client 1 are now in different VLANs and Switch-3 sits on the boundary between the two VLANs. This network is shown in the following figure.

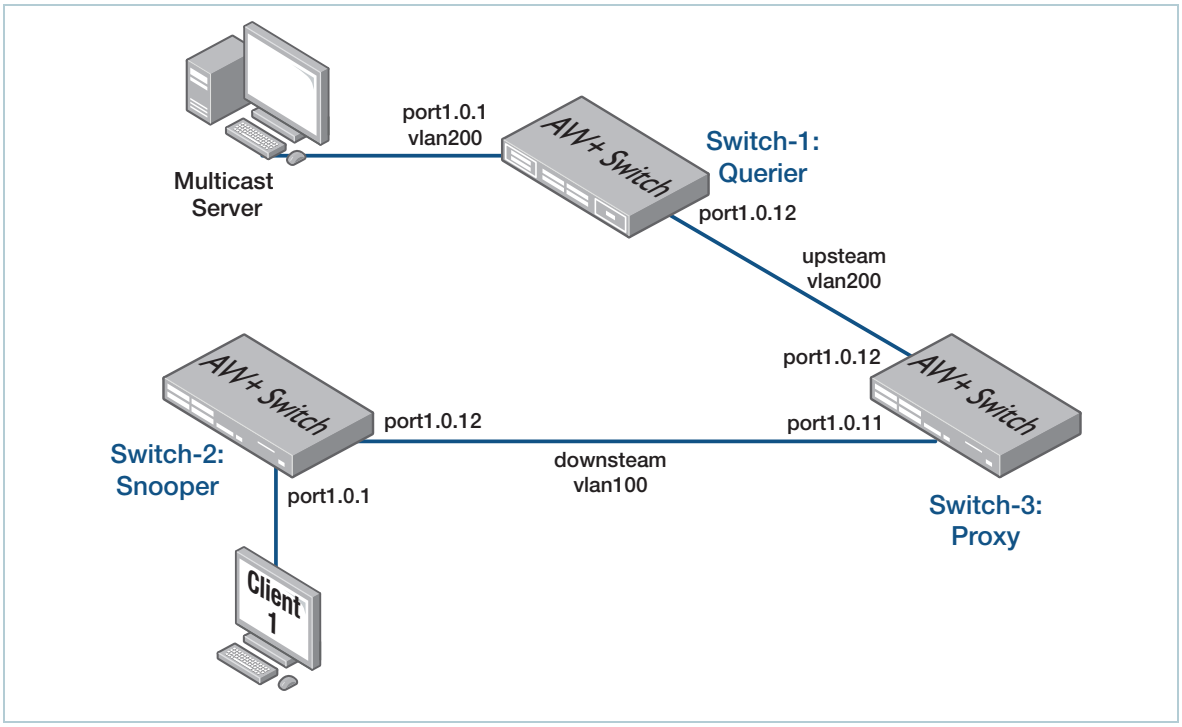

Step 1. **Configure Switch-1**

Switch-1, the closest switch to the multicast source, is an IGMP Querier.

On Switch-1, we have shut down port1.0.11, which is the connection to Switch-2.

```
Switch-1(config)# int port1.0.11
Switch-1(config-if)# shutdown
```

The VLAN has also been changed from vlan100 to vlan200 and we have configured a static route to the vlan100 subnet. So now the configuration for Switch-1 looks as follows:

```
!
hostname Switch-1
!
vlan database
 vlan 200 name vlan200
 !
interface port1.0.1-1.0.10
 switchport access vlan 200
!
interface port1.0.11
 shutdown
  switchport access vlan 200
!
interface port1.0.12-1.0.24
  switchport access vlan 200
!
interface vlan200
 ip address 192.168.200.253/24
 ip igmp
 ip igmp version 2
!
ip route 192.168.100.0/24 192.168.200.254
```

Step 2. **Configure Switch-2**

Switch-2 is an IGMP Snooper. IGMP snooping is enabled by default and does not need any configuration. We have added an IP route to the source's subnet 192.168.200.0/24.

```
!
hostname Switch-2
!
vlan database
 vlan 100 name vlan100
 vlan 100 state enable
!
interface port1.0.1-1.0.24
 switchport access vlan 100
!
interface vlan100
 ip address 192.168.100.252/24
!
ip route 192.168.200.0/24 192.168.100.253
```

Step 3. **Configure Switch-3**

Switch-3 is an IGMP Proxy.

```
hostname Switch-3
!
ip multicast-routing
!
vlan database
 vlan 100 name vlan100
 vlan 200 name vlan200
!
interface port1.0.1-1.0.23
 switchport access vlan 100
!
interface port1.0.24
 switchport access vlan 200
!
!
interface vlan100
 ip address 192.168.100.254/24
 ip igmp
 ip igmp mroute-proxy vlan200
 ip igmp version 2
!
interface vlan200
 ip address 192.168.200.254/24
 ip igmp
 ip igmp proxy-service
 ip igmp version 2
```

Switch-3 has had the most configuration changes in this setup. The upstream interface is vlan200, as this interface is closest to the multicast source, connected to Switch-1. The downstream interface is vlan100, which has the clients requesting the stream 224.12.13.14. Since we are now routing the multicast group 224.12.13.14 from vlan200 to vlan100, we have enabled IP multicast-routing on Switch-3:

```
Switch-3(config)# ip multicast-routing
```

Let us now look in a bit more detail at the meaning of the IGMP Proxy commands. (Explanations below.)

```
Switch-3(config)# int vlan200
Switch-3(config-if)# ip igmp proxy-service
```

On vlan200, we have enabled the IGMP proxy service. What this means is that all associated downstream IGMP mroute proxy interfaces will have their IGMP packets forwarded directly to this interface. The proxy will then transmit these packets towards the server attached to this interface.

```
Switch-3(config)# int vlan100
Switch-3(config-if)# ip igmp mroute-proxy vlan200
```

This downstream mroute proxy interface listens for IGMP Report and Leave messages, and forwards them to the upstream IGMP proxy service interface.

Explanation of IGMP Proxy

When there are no group members

The multicast server streams group 224.12.13.14 to Switch-1 through port 1.0.1. IGMP snooping detects the stream, as you can see by using the command **show ip igmp snooping statistics interface vlan200** on Switch-1.

```
Switch-1#sh ip igmp snooping statistics interface vlan200

IGMP Snooping statistics for vlan200
Group Type:     UnReg MC Group
Interface:      vlan200
Group:          224.12.13.14
Uptime:         00:00:02
Group mode:     Include (Expires: 00:04:17)
Last reporter:  192.168.200.100
```

When a client joins a group

Client 1 (attached to Switch-2, the Snooper) sends an IGMP Membership Report for the group 224.12.13.14. Switch-2 forwards that report message out its Router port, in this case port 1.0.12.

```
Switch-2#sh ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         00:00:59
Group mode:     Include ()
Last reporter:  192.168.100.253

Port member list:
port1.0.12 - 201 secs

Interface:      vlan100
Group:          224.12.13.14
Flags:
Uptime:         00:00:54
Group mode:     Exclude (Expires: 00:03:26)
Last reporter:  192.168.100.10
Source list is empty

Port member list:
port1.0.1 - 206 secs
```

Switch-3—the Proxy—receives the report on its downstream interface, vlan100. Switch-3 then creates a new report with itself as the sender. It sends this report upstream to Switch-1 through vlan200. Output of the commands **show ip igmp groups** and **show ip igmp snooping interface vlan100** show that Switch-3 knows of a client interested in the group 224.12.13.14 through port 1.0.11 on vlan100.

```
Switch-3#sh ip igmp groups

IGMP Connected Group Membership
Group Address    Interface     Uptime    Expires    Last Reporter
224.12.13.14     vlan100       00:58:48  00:02:20   192.168.100.252
224.12.13.14     vlan200       00:58:48  00:04:14   192.168.100.252
```

The group 224.12.13.14 has been registered on both VLANs.

```
Switch-3#sh ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Interface:      vlan100
Group:          224.12.13.14
Flags:
Uptime:         00:02:40
Group mode:     Exclude (Expires: 00:04:09)
Last reporter:  192.168.100.252
Source list is empty

Port member list:
port1.0.11 - 0 secs
```

The downstream client is reached through port1.0.11 via Switch-2. The Last Reporter address is 192.168.100.252, which is the IP address of Switch-2, because Switch-2 applied its IP address to the IGMP reports when it forwarded them.

```
Switch-3#sh ip igmp snooping statistics interface vlan200

IGMP Snooping statistics for vlan200
Group Type:     Router Port Learnt
Interface:      vlan200
Group:          224.0.0.2
Uptime:         00:01:01
Group mode:     Include ()
Last reporter:  192.168.200.253

Port member list:
port1.0.12 - 199 secs


Interface:      vlan200
Group:          224.12.13.14
Flags:
Uptime:         00:03:02
Group mode:     Exclude (Expires: 00:01:28)
Last reporter:  192.168.100.252
Source list is empty

Port member list:
```

On vlan200, the IGMP packets are transmitted out port 1.0.12.

```
Switch-3#sh ip igmp proxy  group

IGMP Connected Proxy Group Membership
Group Address    Interface            Member state
224.12.13.14       vlan200              Delay
```

We can see that the group has been received on vlan200 on the IGMP proxy switch. 'Delay' means that none of the group or source group query timers run for the specified group. The output above does not indicate whether the proxy works or not, it just tells us which groups are registered in the system.

```
Switch-3#sh ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Enabled, Active, Querier, Configured for version 2
 IGMP mroute-proxy interface is vlan200
 Internet address is 192.168.100.253
```

The output above shows that Switch-3 is the active Querier on vlan100 and also that vlan200 is the mroute-proxy interface.

The output below shows that vlan200 has the proxy configured and, importantly, that this interface is not the active Querier.

```
Switch-3#sh ip igmp int vlan200

Interface vlan200 (Index 500)
 IGMP Enabled, Active, Non-Querier, Configured for version 2 proxy-service
 IGMP host version 2
 Internet address is 192.168.200.253
 IGMP interface has 28 group-record states
 IGMP activity: 0 joins, 0 leaves
 IGMP querying router is 192.168.200.254
```

Switch-1 receives the proxied report from Switch-3. Switch-1 notes that Switch-3 is interested in the group 224.12.13.14 and sends the group multicast to Switch-3 on port 1.0.12. Output of the command **show ip igmp snooping statistics interface vlan200** shows the membership that Switch-1 is aware of.

```
Switch-1#sh ip igmp snooping statistics interface vlan200

IGMP Snooping statistics for vlan200
Interface:      vlan200
Group:          224.12.13.14
Uptime:         00:22:31
Group mode:     Exclude (Expires: 00:02:48)
Last reporter:  192.168.200.254
Source list is empty

Port member list:
port1.0.12 - 168 secs
```

When a client leaves a group

When the client on Switch-2 wants to stop receiving the group's multicast stream, it sends an IGMP Leave message. The switches use the above process to transfer the message to Switch-1.

Note the following points about how IGMP Proxy deals with Leave messages:

- The Proxy sends an IGMP Leave Group message via its upstream interface only when the last interface on the Proxy leaves the group.
- The Proxy does not respond to IGMP Join or Leave Group messages received via its upstream interface, but only to those received via downstream interfaces.
- The Proxy does respond to IGMP query messages received via its upstream interface. When the Proxy—Switch-3 in this example—sends a Leave Group message upstream, the upstream IGMP Querier—Switch-1—sends a membership query. Switch-3 takes that query and proxies it to the downstream interface, vlan100, with its own IP address as the source (192.168.100.253). This means any other interested clients on Switch-2 can declare their interest in continuing to receive the multicast stream.

Multiple proxies

A significant feature of AlliedWare Plus is that it can support multiple IGMP proxies on the same switch. An example configuration of multiple IGMP proxies follows:

```
vlan database
 vlan 40 name vlan40
 vlan 50 name vlan50
 vlan 60 name vlan60
 vlan 100 name vlan100
 vlan 200 name vlan200
!
interface port1.0.1-1.0.2
 switchport access vlan 100
!
interface port1.0.3-1.0.4
 switchport access vlan 40
!
interface port1.0.5-1.0.6
 switchport access vlan 50
!
interface port1.0.7-1.0.8
 switchport access vlan 60
!
interface port1.0.24
 switchport access vlan 200
!
interface vlan40
 ip address 192.168.40.254/24
 ip igmp
 ip igmp mroute-proxy vlan50
 ip igmp version 2
!
interface vlan50
 ip address 192.168.50.254/24
 ip igmp
 ip igmp proxy-service
 ip igmp version 2
!
interface vlan60
 ip address 192.168.60.254/24
 ip igmp
 ip igmp mroute-proxy vlan50
 ip igmp version 2
!
interface vlan100
 ip address 192.168.100.254/24
 ip igmp
 ip igmp mroute-proxy vlan200
 ip igmp version 2
!
interface vlan200
 ip address 192.168.200.254/24
 ip igmp
 ip igmp proxy-service
 ip igmp version 2
```

The configuration above has the following interfaces setup as the Proxy Service:

- vlan200
- vlan50

and the following interfaces as the host side (mroute proxy):

- vlan40
- vlan60
- vlan100

Interfaces vlan40 and vlan60 are proxying to vlan50. Quite separately from these, interface vlan100 is proxying to vlan200.

Query Solicitation - Rapid Recovery From Topology Changes

Query Solicitation minimizes loss of multicast data after a topology change. It is a built-in feature of AlliedWare Plus managed Layer 3 switches when running EPSR or spanning tree (STP, RSTP, or MSTP) for loop protection.

Without Query Solicitation, when the underlying link layer topology changes, multicast data flow can stop for up to several minutes, depending on which port goes down and how much of the timeout period was left (see "Why convergence takes so long without Query Solicitation" on page 37). Query Solicitation greatly reduces this disruption.

Note: AlliedWare Plus only supports Query Solicitation for IGMP, and not for MLD.

Query Solicitation operates without configuration in networks of Allied Telesis managed Layer 3 switches running STP, RSTP, MSTP or EPSR. You may find it helpful to manually enable it in the following other situations:

- loop-free networks running IGMP (see "Speeding up IGMP convergence in a non-looped topology" on page 42)
- networks in which not all switches support Query Solicitation (see "Enabling Query Solicitation on multiple switches in a looped topology" on page 43)

How Query Solicitation works

Query Solicitation monitors STP, RSTP, MSTP and EPSR messages for topology changes. When it detects a change, it generates a special IGMP Leave message called a Query Solicit. The switch floods the Query Solicit message to all ports in every VLAN that Query Solicitation is enabled on. When the Querier receives the Query Solicit message, it sends out a General Query and waits for clients to respond with Membership Reports. These Reports update the snooping information throughout the network.

Query Solicit messages have a group address of 0.0.0.0.

Query Solicitation works by default (without you enabling it) on all VLANs on the root bridge in an STP instance and on all data VLANs on the master node in an EPSR instance. By default, the root bridge or master node always sends a Query Solicit message when any of the following events occur:

- an STP BPDU packet with the Topology Change (TC) flag arrives at the root bridge
- an STP port on a switch goes from a Discarding to Forwarding state
- the FDB gets flushed by EPSR

If necessary, you can make clients respond more quickly to the General Query by tuning the IGMP timers, especially the "Max Query Response Interval" on page 78.

The following figure shows how Query Solicitation works when a port goes down:

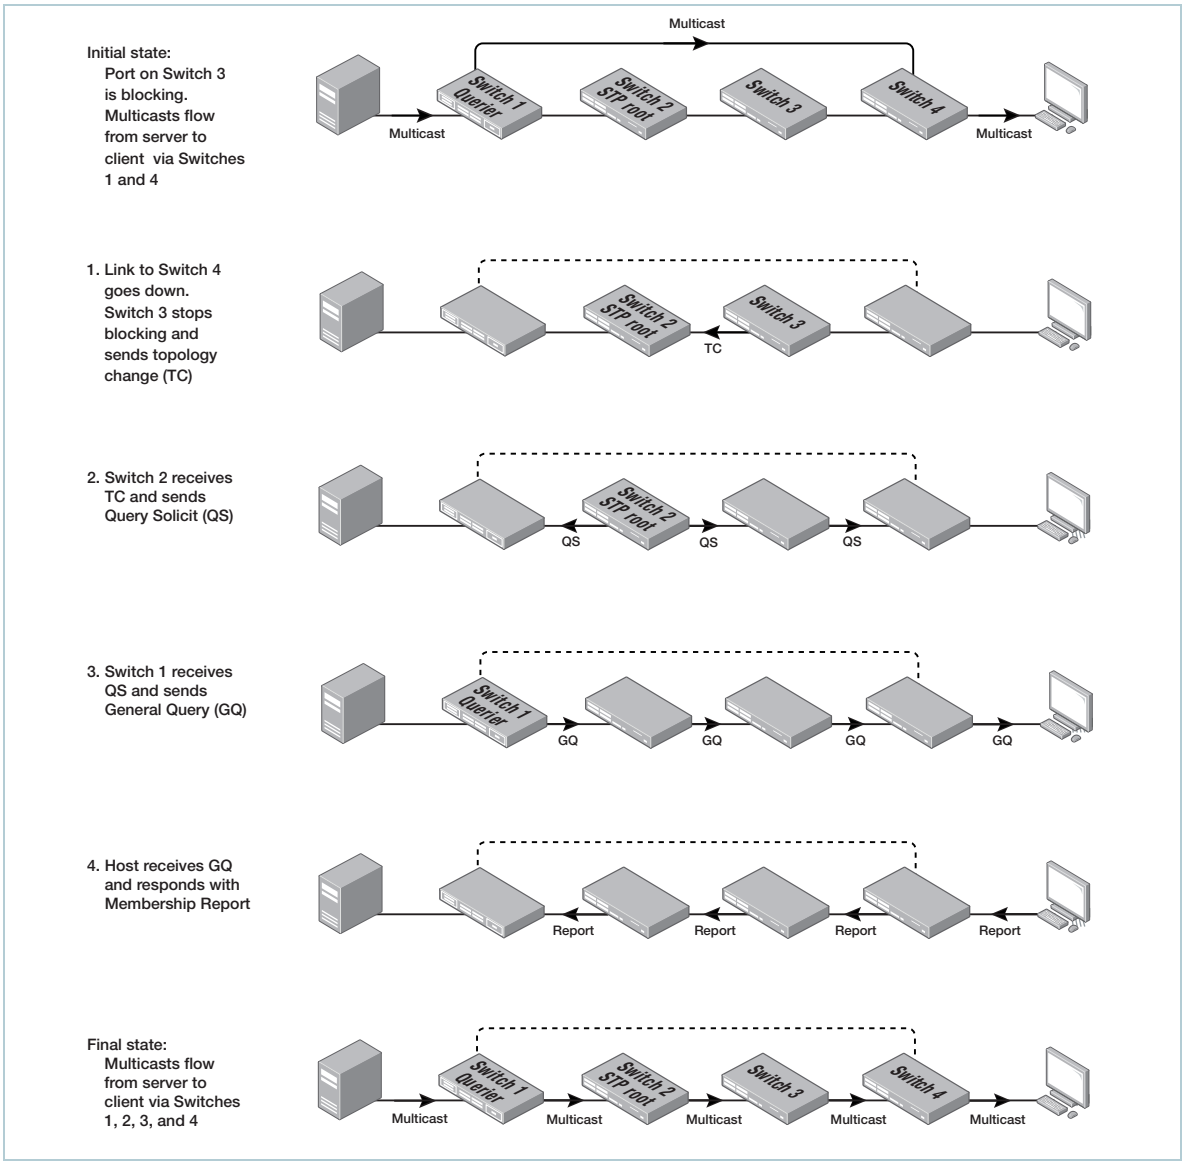

Why convergence takes so long without Query Solicitation

This section illustrates IGMP convergence in a simple network that does not need STP because it has no switch loops. In this case we have disabled STP on all switches. Query Solicitation is disabled by default in networks like this, because no switch is an STP root bridge or an EPSR master node.

In this network, it takes up to 125 seconds for multicasting to recover after a port comes back up. This section explains the reason for the slow convergence. "Speeding up IGMP convergence in a non-looped topology" on page 42 explains the solution.

Example

The following figure shows the network for the example in this section:

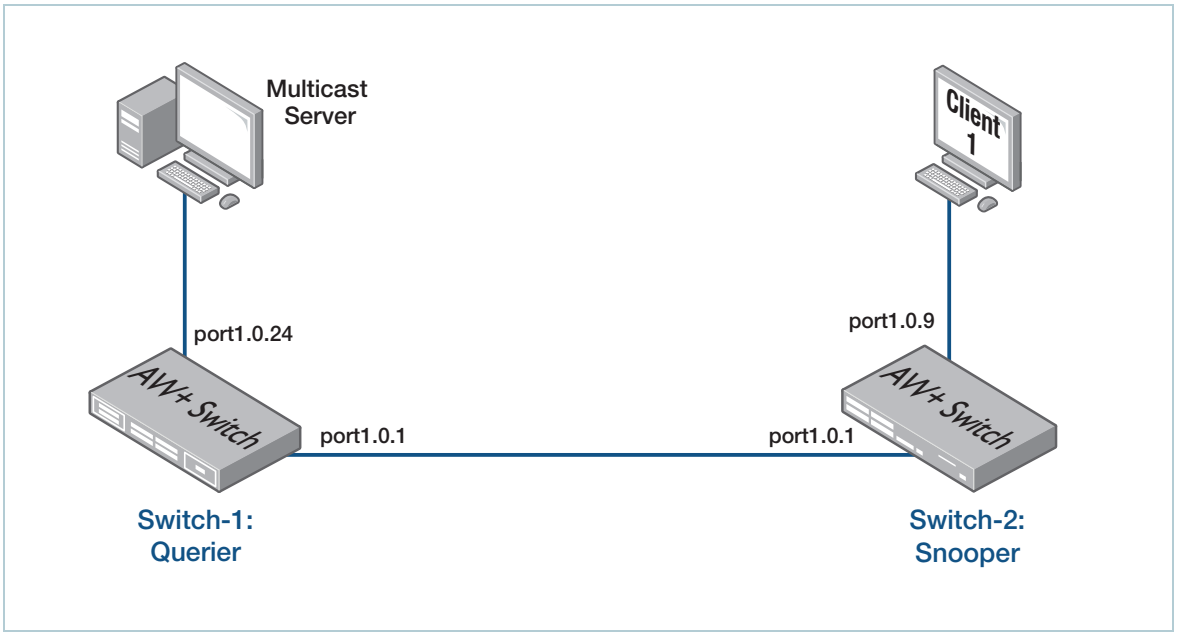

The example considers what happens when a port comes up. When the port was down, the client stopped receiving multicasts because there was no backup route available. The example shows how the network recovers. The multicast group is 224.12.13.14.

Step 1. **Configure Switch-1**

Switch-1 is configured with IGMP, which makes it the IGMP Querier in this network.

```
!
hostname Switch-1
!
no spanning-tree rstp enable
!
interface vlan1
 ip address 192.168.1.1/24
 ip igmp
```

Step 2. **Configure Switch-2**

Switch-2 is an IGMP Snooper. It forwards multicast packets and IGMP messages as required. IGMP snooping is enabled by default and does not need any configuration.

```
!
hostname Switch-2
!
!no spanning-tree rstp enable
!
interface vlan1
 ip address 192.168.1.2/24
```

Explanation from the perspective of Switch-2, the Snooper

When link is up

When the link is connected (all ports are up), the Snooper has entries for two ports:

- port 1.0.9, which is the Snooper's connection to the client. The Snooper sends the multicast stream out this port, as well as forwarding Queries (the Snooper floods Queries out all its ports).
- port 1.0.1, which is the Snooper's connection to the Querier. This is a Router Port entry, so the Snooper forwards Reports out this port. The output of the command **show ip igmp snooping statistics interface vlan1** shows both entries.

```
Switch-2#sh ip igmp snooping statistics interface vlan1

IGMP Snooping statistics for vlan1
Group Type:     Router Port Learnt
Interface:      vlan1
Group:          224.0.0.2
Uptime:         00:03:18
Group mode:     Include ()
Last reporter:  192.168.1.1

Port member list:
port1.0.1 - 218 secs

Interface:      vlan1
Group:          224.12.13.14
Flags:
Uptime:         00:01:49
Group mode:     Exclude (Expires: 00:03:41)
Last reporter:  192.168.1.10
Source list is empty

Port member list:
port1.0.9 - 221 secs
```

When link goes down

When we disconnect port 1.0.1 on the Snooper, the Router Port disappears, and the entry for 224.12.13.14 is still present until it times out.

```
Switch-2#sh ip igmp snooping statistics interface vlan1

IGMP Snooping statistics for vlan1
Interface:      vlan1
Group:          224.12.13.14
Flags:
Uptime:         00:02:56
Group mode:     Exclude (Expires: 00:02:33)
Last reporter:  192.168.1.10
Source list is empty

Port member list:
port1.0.9 - 154 secs
```

Once the link is reconnected, we see that the Router Port has not been learnt.

```
Switch-2#sh ip igmp snooping statistics interface vlan1

IGMP Snooping statistics for vlan1
Interface:      vlan1
Group:          224.12.13.14
Flags:
Uptime:         00:05:01
Group mode:     Exclude (Expires: 00:00:28)
Last reporter:  192.168.1.10
Source list is empty

Port member list:
port1.0.9 - 29 secs
```

Eventually, the Querier sends an IGMP Query, which the Snooper receives on port 1.0.1. This restores the All Groups port on the Snooper. By default the Querier sends General Queries every 125 seconds, so the IGMP convergence delay will be up to 125 seconds with the default settings. For more information about this timeout, see "Configurable IGMP/MLD Timers and Counters" on page 71—but do not change the timeout without very carefully considering the effect on your network.

When the Snooper receives the General Query, it forwards it out all its ports. The client responds with a Report. The Snooper forwards the Report out its All Groups port towards the Querier. The Querier responds by sending the multicast stream to the Snooper, which forwards the multicast stream out port 1.0.9 to the client.

```
Switch-2#sh ip igmp snooping statistics interface vlan1

IGMP Snooping statistics for vlan1
Group Type:     Router Port Learnt
Interface:      vlan1
Group:          224.0.0.2
Uptime:         00:01:15
Group mode:     Include ()
Last reporter:  192.168.1.1

Port member list:
port1.0.1 - 216 secs

Interface:      vlan1
Group:          224.12.13.14
Flags:
Uptime:         00:06:32
Group mode:     Exclude (Expires: 00:03:36)
Last reporter:  192.168.1.10
Source list is empty

Port member list:
port1.0.9 - 217 secs
```

Explanation from the perspective of Switch-1, the Querier

When link is up

When the link is connected (all ports are up), the Querier has an entry for port 1.0.1, so it sends the group 224.12.13.14 out port 1.0.1. The output of the command **show ip igmp snooping statistics interface vlan1** shows this entry.

```
Switch-1#show ip igmp snooping statistics interface vlan1

IGMP Snooping statistics for vlan1
Interface:      vlan1
Group:          224.12.13.14
Uptime:         16:49:17
Group mode:     Exclude (Expires: 00:03:18)
Last reporter:  192.168.1.2
Source list is empty

Port member list:
port1.0.1 - 198 secs
```

When link goes down

When we disconnect port 1.0.1 on the Snooper, the port disappears. The Querier is still receiving the multicast stream from the server, so the group entry remains. We can see that the group 224.12.13.14 is now an ‘UnReg MC Group’.

```
Switch-1#sh ip igmp snooping statistics interface vlan1

IGMP Snooping statistics for vlan1
Group Type:     UnReg MC Group
Interface:      vlan1
Group:          224.12.13.14
Uptime:         00:00:05
Group mode:     Include (Expires: 00:04:15)
Last reporter:  192.168.1.5
```

When link comes up again

When we reconnect port 1.0.1 on the Snooper, the port does not reappear in the IGMP output because the Querier has not yet received a Report over it. Therefore, the Querier does not start forwarding the multicast stream out the port at this time.

Eventually, the Querier sends an IGMP Query out all its ports. In response, it receives a Report from the client (via the Snooper). This restores the port entry and the Querier starts sending the multicast stream again.

The output of the commands **show ip igmp snooping statistics interface vlan1** and **show ip igmp groups** both show this restored entry.

```
Switch-1#sh ip igmp snooping statistics interface vlan1

IGMP Snooping statistics for vlan1
Interface:      vlan1
Group:          224.12.13.14
Uptime:         00:00:12
Group mode:     Exclude (Expires: 00:04:08)
Last reporter:  192.168.1.2
Source list is empty

Port member list:
port1.0.1 - 248 secs
```

```
Switch-1#sh ip igmp groups

IGMP Connected Group Membership
Group Address    Interface            Uptime   Expires  Last Reporter
224.12.13.14     vlan1                00:00:21 00:03:59 192.168.1.2
```

Speeding up IGMP convergence in a non-looped topology

The previous section described how it can take up to 125 seconds for multicasting to recover in a non-looped topology after a port comes back up. You can speed up convergence simply by enabling RSTP. This enables the network to use Query Solicitation and means that multicasting resumes within 3 seconds of the link coming up.

Even though there is no loop in the network, one of the switches becomes the STP root bridge—it does not matter which switch does this. When the link comes up, the root bridge detects the topology change and sends a Query Solicitation.

We disabled RSTP in the previous example. Even though RSTP is enabled by default in AlliedWare Plus, we have to re-enable it for this example:

```
Switch-1# confifure terminal
Switch-1(config)# spanning-tree rstp enable
```

Enabling Query Solicitation on multiple switches in a looped topology

On networks that use spanning tree or EPSR, Query Solicitation is not normally required on switches other than the STP root bridge or EPSR master node. Therefore, it is only enabled by default on the root bridge and the master node.

However, in some networks you may need to turn on Query Solicitation on all switches—for example, if the network includes other switches that do not support Query Solicitation and therefore the STP root bridge may be a switch that does not send Query Solicit messages. To enable Query Solicitation, use the commands:

```
awplus# configure terminal
awplus(config)# ip igmp snooping tcn query solicit
```

Every switch that has Query Solicitation enabled sends a Query Solicit message when it detects a topology change. Enabling it on multiple switches means you get multiple messages, but has no other disadvantage.

The following figure shows the packet flow for a four-switch network with Query Solicitation enabled on all the switches.

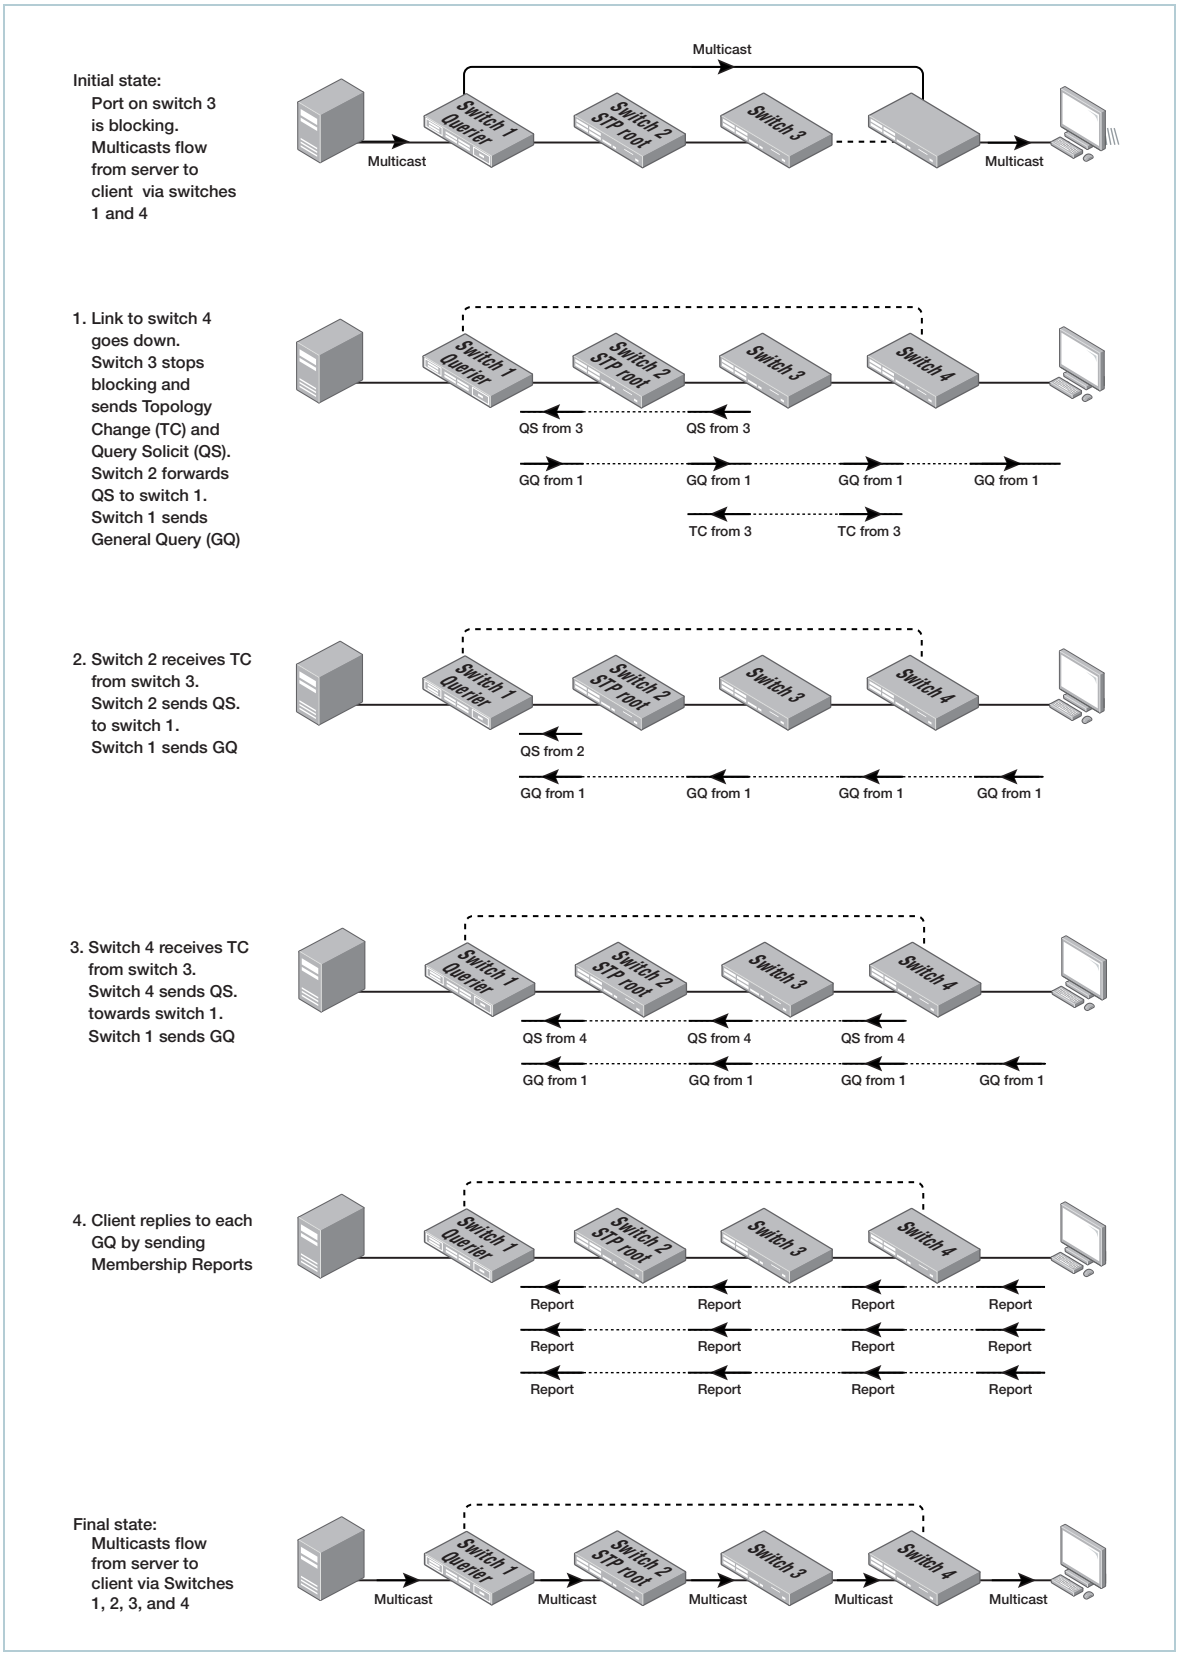

One topology change caused three Query Solicits, three General Queries, and three Reports.

IGMP/MLD Filtering (controlling multicast distribution)

IGMP/MLD filtering lets you control the distribution of multicast services on each switch port. Filtering is useful for subscription services when clients must be explicitly authorized to view a multicast stream.

Example

This example shows how to stop a host joining the groups 224.0.1.22 and FF02::1:1000 and allow it to join all other groups. It uses the same network configuration as "IGMP/MLD Snooping" on page 10. For convenience, the diagram is reproduced below.

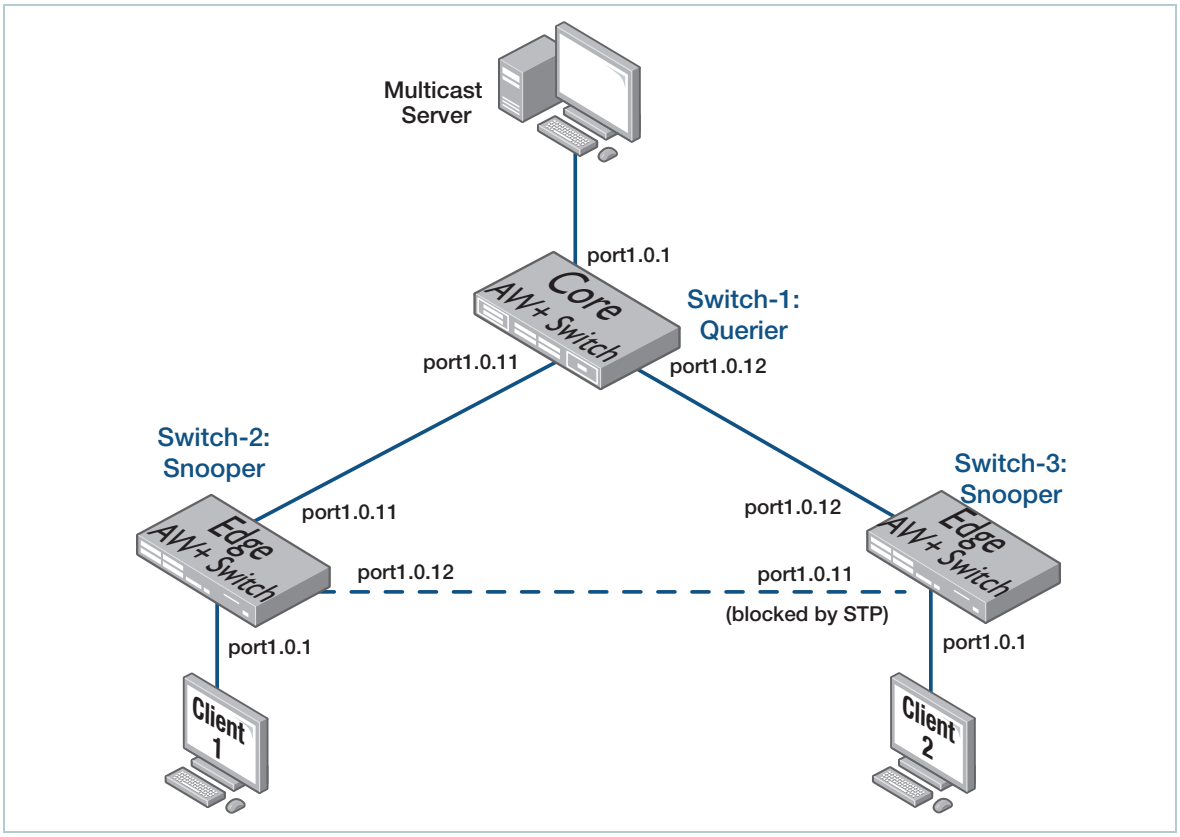

The network contains a Windows 2000 workstation that regularly sends SVRLOC messages (an IGMP Membership Report for 224.0.1.22) and MLD membership report for FF02::1:1000. These groups gets added to the list of groups in vlan100 on Switch-1, as shown in the output of the following **show** commands.

```
Switch-1>show ip igmp groups

IGMP Connected Group Membership
Group Address    Interface            Uptime   Expires  Last Reporter
224.0.1.22       vlan100              00:03:27 00:03:04 0.0.0.0
```

```
Switch-1>show ipv6 MLD groups

MLD Connected Group Membership
Group Address           Interface          Uptime           Expires
Last Reporter
FF02::1:1000          vlan100          00:03:27        00:03:04
fe80::eecd:6dff:fe6b:4783
```

We do not need to receive this multicast, so we will filter it out.

Step 1. **Configure Switch-1**

Switch-1, the closest switch to the multicast source, is an IGMP/MLD Querier. AlliedWare Plus uses access lists to accomplish the desired effect. Add an Access Control List (ACL) to a VLAN interface configured for IGMP/MLD, IGMP/MLD Snooping or IGMP Proxy:

```
Switch-1(config)# ip igmp access-group
Switch-1(config)# ipv6 mld access-group
```

The access control list is used to control and filter the multicast groups learnt on the VLAN interface.

To configure an ACL to block 224.0.1.22 and to attach the ACL to vlan100, use these commands:

```
Switch-1#configure terminal
Switch-1(config)#access-list 1 deny 224.0.1.22 0.0.0.0
Switch-1(config)#int vlan100
Switch-1(config-if)#ip igmp access-group 1
```

```
Switch-1#show access-list 1

Standard IP access list 1
   10 deny   224.0.1.22
```

To configure an ACL to block FF02::1:1000 and to attach the ACL to vlan100, use these commands:

```
Switch-1#configure terminal
Switch-1(config)#ipv6 access-list standard no-svrloc deny  FF02::1:1000/
                 128
Switch-1(config)#int vlan100
Switch-1(config-if)#ipv6 mld access-group no-svrloc
```

Switch-1 configuration:

```
!
hostname Switch-1
!
access-list 1 deny 224.0.1.22
!
ipv6 access-list standard no-svrloc deny FF02::1:1000/128
spanning-tree mode rstp
ipv6 forwarding
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
!
interface vlan100
 ip address 192.168.100.254/24
 ip igmp
 ip igmp access-group 1
 ip igmp version 2
 ipv6 address 2001:0db8::12:251/64
 ipv6 enable
 ipv6 mld
 ipv6 mld access-group no-svrloc
 ipv6 mld version 1
```

Switch-2 and Switch-3's configuration do not change from the first example ("Configuration example" on page 12).

We look at the command **show ip igmp snooping statistics interface vlan100**, and since applying the filter, the stream 224.0.1.22 is no longer in the output:

```
Switch-1#sh ip igmp snooping statistics interface vlan100
IGMP Snooping statistics for vlan100
```

IGMP/MLD Throttling

IGMP/MLD throttling allows you to limit the number of streams that subscribers may access at a given time, for example, to protect from bandwidth over subscription. When the number of multicast group memberships associated with a VLAN or (for IGMP) switch port reaches the configured limit, it can either deny further Membership Reports, or replace an existing membership with the new group.

IGMP/MLD filtering and throttling can be applied separately or together. The switch applies the filters first, then subjects any multicast group memberships passed by the filter to the limits imposed by throttling.

Example of per-VLAN throttling

This example builds on "IGMP/MLD Filtering (controlling multicast distribution)" on page 45 and uses the same network configuration as "IGMP/MLD Snooping" on page 10. For convenience, the diagram is reproduced below.

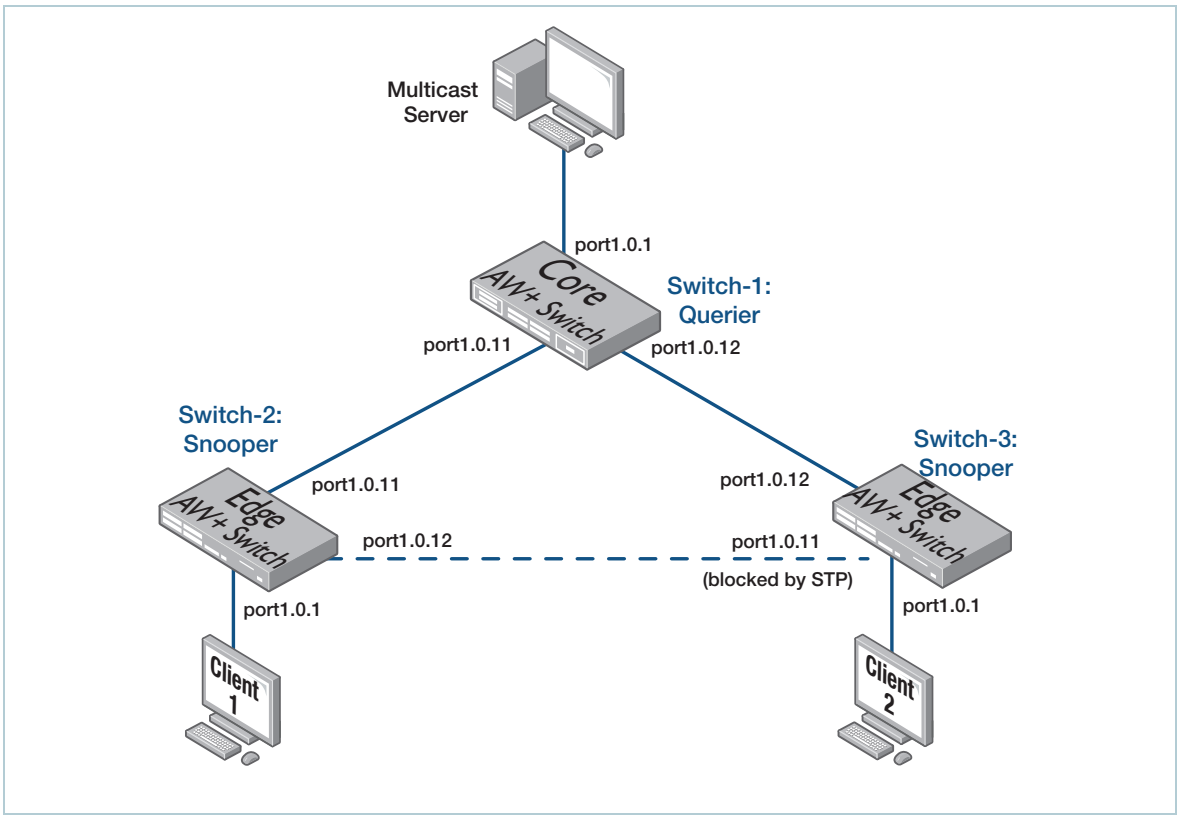

Step 1. **Configure Switch-1**

Switch-1 is an IGMP/MLD Querier. It has the same filter configured as in the previous example, to deny the multicast groups 224.0.1.22.

```
!
hostname Switch-1
!
access-list 1 deny 224.0.1.22
!
Switch-1(config)#ipv6 access-list standard no-svrloc deny FF02::1:1000/128
!
spanning-tree mode rstp
ipv6 forwarding
vlan database
 vlan 100 state enable
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
!
interface vlan100
 ip address 192.168.100.254/24
 ip igmp
 ip igmp access-group 1
 ip igmp version 2
 ipv6 address 2001:0db8::12:251/64
 ipv6 enable
 ipv6 mld
 ipv6 mld access-group no-svrloc
 ipv6 mld version 1
```

Step 2. **Configure Switch-2**

Switch-2 is an IGMP/MLD Snooper. IGMP/MLD snooping is enabled by default and does not need any configuration.

Switch-2 is limited to three IPv4 and five IPv6 multicast groups on vlan100. We have also added the filters that were configured on Switch-1 to deny the multicast groups 224.0.1.22. and FF02::1:1000.

```
!
hostname Switch-2
ipv6 forwarding
access-list 1 deny 224.1.0.22
!
ipv6 access-list standard no-svrloc deny FF02::1:1000/128
!
spanning-tree mode rstp
ipv6 forwarding
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
!
interface vlan100
 ip igmp access-group 1
 ip igmp limit 3
 ipv6 mld access-group no-svrloc
 ipv6 mld limit 5
```

Step 3. **Configure Switch-3**

Switch-3 is also an IGMP/MLD Snooper.

```
!
hostname Switch-3
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
  switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
```

Explanation of IGMP/MLD throttling

In this example, Switch-2's configuration limits vlan100 to three concurrent IPv4 multicast groups. Consider Switch-2 after a client on port 1.0.1 has tried to join the three groups from 224.12.13.14 - 224.12.13.16. Output from the command **show ip igmp snooping statistics interface vlan100** shows the two memberships.

Table 4: Output from the show ip igmp snooping statistics interface vlan100 command

```
Switch-2#show ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         02:12:43
Group mode:     Include ()
Last reporter:  192.168.100.254

Port member list:
port1.0.11 - 172 secs

Interface:      vlan100
Group:          224.12.13.14
Flags:
Uptime:         00:51:03
Group mode:     Exclude (Expires: 00:02:54)

Last reporter:  192.168.100.10
Source list is empty

Port member list:
port1.0.1 - 175 secs

Interface:      vlan100
Group:          224.12.13.15
Flags:
Uptime:         00:16:00
Group mode:     Exclude (Expires: 00:02:56)
Last reporter:  192.168.100.10
Source list is empty

Port member list:
port1.0.1 - 177 secs

Group Type:     UnReg MC Group
Interface:      vlan100
Group:          239.255.255.253
Uptime:         00:05:14
Group mode:     Include (Expires: 00:01:26)
Last reporter:  192.168.1.100
```

The switch has three entries in the IP IGMP snooping table, and even though the client attached to Switch-2 was trying to access the streams 224.12.13.14, 224.12.13.15 and 224.12.13.16, it has only joined streams 224.12.13.14 and 224.12.13.15, as Switch-2 has also got an entry for the 239.255.255.253 group (SLP multicast address). The limit of three has been reached, so 224.12.13.16 is not present in the table and the client cannot receive this group.

Per-port throttling of IGMP groups

As well as limiting per VLAN, as in the above example, the throttling can also be applied on a per-port basis from Version 5.4.6-1.x onwards for IGMP. You can simply set a limit, per switch port, on the number of IGMP groups that clients can join. This stops a single client from using all the switch's available group-entry resources, and ensures that clients on all ports have a chance to join IGMP groups.

To set the limit, go into interface mode for the switch port or ports and use the command:

```
awplus(config-if)# ip igmp maximum-groups <0-65535>
```

The default is 0, which means no limit.

We recommend using this with IGMP snooping fast leave on the relevant VLANs. To enable fast leave, use the command:

```
awplus(config-if)# ip igmp snooping fast-leave
```

The device keeps count of the number of groups learned by each port. This counter is incremented when group joins are received via IGMP reports. It is decremented when:

- Group leaves are received via leave messages or reports
- Group memberships time out

Also, the port's group counter is cleared when:

- The port goes down
- You run the command **clear igmp groups ***
- The port is removed from a VLAN
- The port is on a VCStack back-up member, and that member reboots or otherwise leaves the stack.

You can see the current value of the group counter by using either of the commands:

```
awplus# show ip igmp snooping statistics interface <port-list>
awplus# show ip igmp interface <port>
```

For example, to display information about port1.0.3, use either of the following commands:

```
awplus# show ip igmp snooping statistics interface port1.0.3
```

```
IGMP Snooping statistics for port1.0.3
Maximum groups limit set: 10
Number of groups port belongs to: 2
```

```
awplus# show ip igmp interface port1.0.3
```

```
IGMP information for port1.0.3
Maximum groups limit set: 10
Number of groups port belongs to: 2
```

Static IGMP/MLD

Static IGMP/MLD enables you to configure a switch with specified group-to-interface or group-to-port mappings, which you may want to do if:

- your network includes hosts that cannot send IGMP/MLD Membership Reports
- you need to guarantee that a specific multicast stream is instantly available on a port, without any delay from the joining process

A common usage of Static IGMP/MLD is for protocols like Service Location Protocol (SLP). This protocol sends out multicast packets that need to be forwarded to designated ports. You may want SLP packets to be forwarded to ports that have servers who need to respond to these packets. Static IGMP/MLD allows you to specify that traffic for this group should go to hosts who will respond to, or are interested in, these messages.

Example

In this example, we will start by setting an IGMP static entry for the group 224.12.13.14 to go to port 1.0.5 on Switch-1 (part of vlan100 and a static MLD entry for the group FF02::1:1000 to go to the same port). On that port, we have attached a host that has no multicast client software running. After examining the effect of static IGMP/MLD on Switch-1, we will add a static IGMP entry and a static MLD entry to Switch-3 and consider the effect this has on multicasting through the network. The network for this example has three switches in a loop and is shown in the following figure.

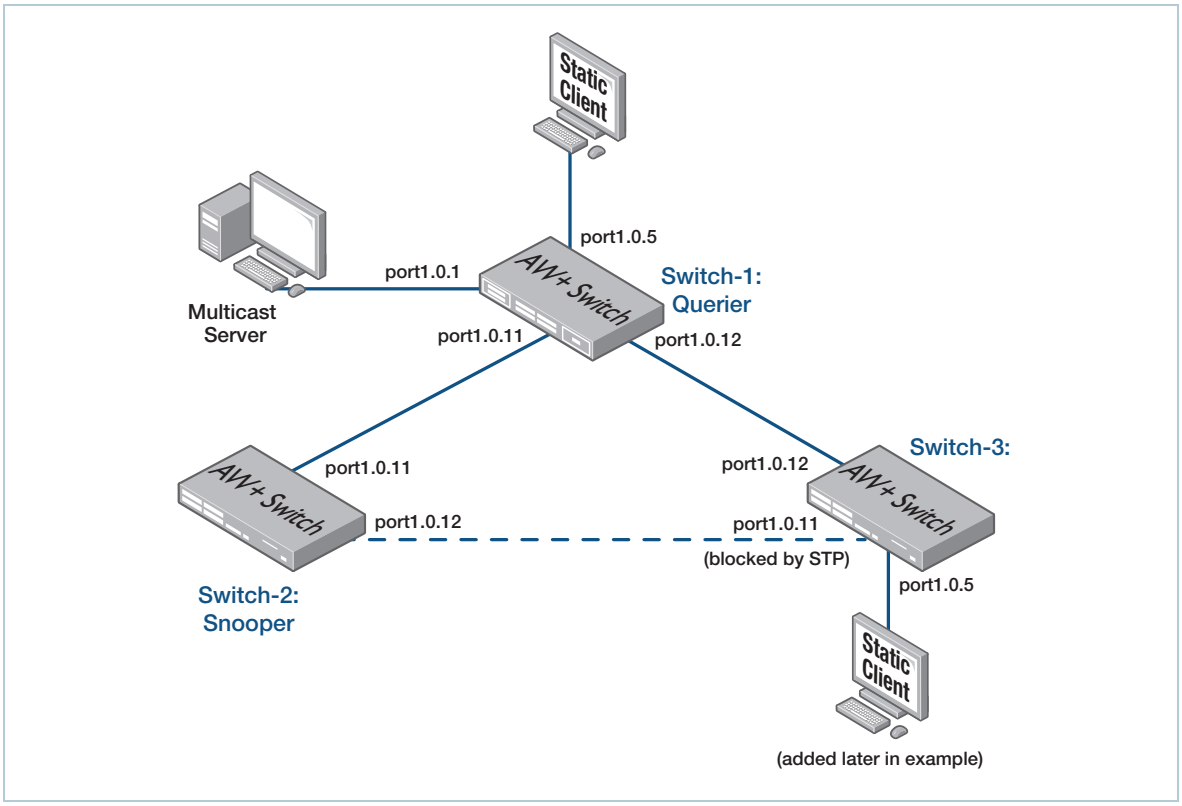

Step 1. **Configure Switch-1**

- Switch-1 is an IGMP/MLD Querier and has the static IGMP and MLD entries. Static IGMP/MLD also requires you to:
- add an IP address to the interface to which you will attach the static entry
- enable IGMP and MLD
- enable the interface as an IGMP/MLD interface

The commands for adding the static entries to the switch are:

```
Switch-1# configure terminal
Switch-1(config)# int vlan100
Switch-1(config-if)# ip igmp static-group 224.12.13.14 int port1.0.5
Switch-1(config-if)# ipv6 mld static-group FF02::1:1000 int port1.0.5
```

Configuration of Switch-1

```
!
hostname Switch-1
!
spanning-tree mode rstp
ipv6 forwarding
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
!
interface vlan100
 ip address 192.168.100.254/24
 ip igmp
 ip igmp static-group 224.12.13.14 interface port1.0.5
 ip igmp version 2
 ipv6 address 2001:0db8::12:252/64
 ipv6 enable
 ipv6 mld
 ipv6 mld version 1
 ipv6 mld static-group FF02::1:1000 int port1.0.5
```

Step 2. **Configure Switch-2**

Switch-2 is an IGMP Snooper.

```
!
hostname Switch-2
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
```

Step 3. **Configure Switch-3**

Switch-3 is also an IGMP Snooper.

```
!
hostname Switch-3
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
  switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
!
```

Explanation of static IGMP/MLD

When the IGMP static entry is created on Switch-1, entries immediately appear in the IGMP snooping table and the IGMP table.

```
Switch-1#show ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Interface:      vlan100
Group:          224.12.13.14
Uptime:         00:18:08
Group mode:     Exclude (Static)
Last reporter:  0.0.0.0
Source list is empty

Port member list:
port1.0.5 - 0 secs

Switch-1#sh ip igmp groups
IGMP Connected Group Membership
Group Address    Interface            Uptime     Expires  Last Reporter
224.12.13.14     vlan100              00:19:01   stopped  0.0.0.0
```

The reason that it states ‘stopped’ in the ‘Expires’ column is that it is a static entry, so this entry will never time out. If there are no clients that have sent an IGMP Membership Report for the static group then the 'Expires' entry shows 'stopped'. If the switch sees any IGMP Membership Reports for the static group from clients, then the 'Expires' entry shows 'static'.

When the group 224.12.13.14 starts streaming into Switch-1, we can use the command **show interface port1.0.5** to watch the number of multicast packets sent on port1.0.5 increase. This means that the link is up and the static IGMP entry is working.

```
Switch-1#show int port1.0.5

Interface port1.0.5
  Scope: both
  Link is UP, administrative state is UP
  Thrash-limiting
    Status Not Detected, Action learn-disable, Timeout 1(s)
  Hardware is Ethernet, address is 0000.cd29.98ab
  index 5005 metric 1 mru 1522
  current duplex full, current speed 100, current polarity auto
  configured duplex auto, configured speed auto, configured polarity auto
  <UP,BROADCAST,RUNNING,MULTICAST>
  SNMP link-status traps: Disabled
   input packets 541, bytes 47972, dropped 0, multicast packets 147
   output packets 2919, bytes 1031906, multicast packets 2259 broadcast packets 314
  Time since last state change: 0 days 00:51:56
```

```
Switch-1#show int port1.0.5

Interface port1.0.5
  Scope: both
  Link is UP, administrative state is UP
  Thrash-limiting
    Status Not Detected, Action learn-disable, Timeout 1(s)
  Hardware is Ethernet, address is 0000.cd29.98ab
  index 5005 metric 1 mru 1522
  current duplex full, current speed 100, current polarity auto
  configured duplex auto, configured speed auto, configured polarity auto
  <UP,BROADCAST,RUNNING,MULTICAST>
  SNMP link-status traps: Disabled
    input packets 543, bytes 48284, dropped 0, multicast packets 149
    output packets 3518, bytes 1842552, multicast packets 2858 broadcast packets 314
  Time since last state change: 0 days 00:52:03
```

When we add a static entry on another switch

Now we add static IGMP and MLD entries on port1.0.5 of Switch-3, by adding the configuration below:

```
!
interface vlan100
 ip address 192.168.100.253/24
 ip igmp
 ip igmp static-group 224.12.13.14 interface port1.0.5
 ip igmp version 2
 ipv6 address 2001:0db7::12:252/64
 ipv6 enable
 ipv6 mld
 ipv6 mld version 1
 ipv6 mld static-group FF02::1:1000 int port1.0.5
```

Both switches are potential IGMP/MLD Queriers and Switch-3 becomes the IGMP and MLD Querier. This is because we gave Switch-3 lower IP addresses (192.168.100.253 and 2001:0db7::12:252/64) than Switch-1 (192.168.100.254 and 2001:0db8::12:252/64).

To see the effect that the new configuration has on Switch-1, we can check the IGMP snooping and IGMP tables.

The IGMP snooping table shows that Switch-1 now has an All Groups entry because it is no longer the Querier. The IGMP table also shows that Switch-1 is not the Querier.

```
Switch-1#sh ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         00:46:58
Group mode:     Include ()
Last reporter:  192.168.100.253

Port member list:
port1.0.12 - 228 secs

Interface:      vlan100
Group:          224.12.13.14
Uptime:         00:57:06
Group mode:     Exclude (Static)
Last reporter:  0.0.0.0
Source list is empty

Port member list:
port1.0.5 - 0 secs
```

```
Switch-1#sh ip igmp groups

IGMP Connected Group Membership
Group Address    Interface          Uptime    Expires  Last Reporter
224.12.13.14     vlan100            00:57:56  stopped  0.0.0.0
```

We can see the static entry on Switch-3 by checking the IGMP snooping and IGMP tables.

```
Switch-3#sh ip igmp snooping statistics interface vlan100

IGMP Snooping statistics for vlan100
Interface:      vlan100
Group:          224.12.13.14
Flags:         SG
Uptime:         00:50:04
Group mode:     Exclude (Expires: 00:02:42, Static)
Last reporter:  192.168.100.254
Source list is empty

Port member list:
port1.0.5 - 0 secs
```

```
Switch-3#sh ip igmp groups

IGMP Connected Group Membership
Group Address    Interface          Uptime   Expires  Last Reporter
224.12.13.14     vlan100            00:52:16   static 192.168.100.254
```

The server connected to port1.0.1 on Switch-1 is transmitting group 224.12.13.14, so we can see from the first output that Switch-3 has received the group 224.12.13.14 on port 1.0.12 and the static entry is on port 1.0.5.

If the multicast server was attached to Switch-3 instead of Switch-1, we would have to change Switch-3's configuration. We would need to add a static entry for the port that Switch-3 uses to connect to Switch-1 (port 1.0.12). Although this is unnecessary in this scenario, we will do it to demonstrate the effect, by using the following command:

```
awplus(config-if)# ip igmp static-group 224.12.13.14 interface port1.0.12
```

To see the new static entry, we use the commands **show ip igmp snooping statistics interface vlan100** and **show ip igmp group**; to see multicast packets streaming, we use the command **show int port1.0.5** and **show int port1.0.12** to view the counters.

```
Switch-3#sh ip igmp groups

IGMP Connected Group Membership
Group Address    Interface        Uptime   Expires  Last Reporter
224.12.13.14     vlan100          01:12:18   static 192.168.100.254
```

```
Switch-3#sh ip igmp snooping statistics interface vlan100
IGMP Snooping statistics for vlan100
Interface:      vlan100
Group:          224.12.13.14
Flags:         SG
Uptime:         01:12:41
Group mode:     Exclude (Expires: 00:03:23, Static)
Last reporter:  192.168.100.254
Source list is empty

Port member list:
port1.0.5 - 0 secs
port1.0.12 - 0 secs
```

```
Switch-3#sh int port1.0.5

Interface port1.0.5
  Scope: both
  Link is UP, administrative state is UP
  Thrash-limiting
    Status Not Detected, Action learn-disable, Timeout 1(s)
  Hardware is Ethernet, address is 0015.77c2.4d55
  index 5005 metric 1 mru 1500
  current duplex full, current speed 100, current polarity mdix
  configured duplex auto, configured speed auto, configured polarity auto
  <UP,BROADCAST,RUNNING,MULTICAST>
  SNMP link-status traps: Disabled
    input packets 19, bytes 1885, dropped 0, multicast packets 8
    output packets 523, bytes 37273, multicast packets 417 broadcast packets 106
  Time since last state change: 0 days 00:13:29
```

```
Switch-3#sh int port1.0.12
Interface port1.0.12
  Scope: both
  Link is UP, administrative state is UP
  Thrash-limiting
    Status Not Detected, Action learn-disable, Timeout 1(s)
  Hardware is Ethernet, address is 0015.77c2.4d55
  index 5012 metric 1 mru 1500
  current duplex full, current speed 1000, current polarity mdi
  configured duplex auto, configured speed auto, configured polarity auto
  <UP,BROADCAST,RUNNING,MULTICAST>
  SNMP link-status traps: Disabled
    input packets 58783, bytes 10599594, dropped 0, multicast packets 50137
    output packets 1725, bytes 197790, multicast packets 713 broadcast packets 1
012
  Time since last state change: 1 days 00:05:40
```

We can see that the counters have incremented on both ports.

When a static entry's port goes down

Finally, note that when the port attached to a static entry goes down, the static entry remains. You can see from the output of the command **show ip igmp snooping statistics interface vlan100** that port 5 has been disconnected.

```
Switch-1#sh ip igmp snooping statistics interface vlan100
IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         01:35:25
Group mode:     Include ()
Last reporter:  192.168.100.253

Port member list:
port1.0.12 - 227 secs

Interface:      vlan100
Group:          224.12.13.14
Uptime:         01:45:33
Group mode:     Exclude (Static)
Last reporter:  0.0.0.0
Source list is empty

Port member list:
port1.0.5 - 0 secs
```

Static router ports

A Router Port (or All Groups port) is a port that the switch identifies as connecting the path to a routing device. The switch automatically makes ports into Router Ports when it detects incoming packets such as IGMP queries or OSPF messages, because these indicate a connection to a routing device.

The switch sends **all** multicast streams it receives out all Router Ports. This is why Router Ports also called All Groups ports.

If you want to have multicast streams flooded out a port, and that port does not happen to be a Router Port (because it has not received packets like IGMP queries or OSPF messages), then you can configure the port as a static Router Port. To do this, use the command:

```
ip igmp snooping mrouter interface <port>
```

For example, to make ports 1.0.4, 1.0.5 and 1.0.6 flood all multicast packets that arrive on VLAN20, then use the following commands to configure them as static Router Ports:

```
awplus# configure terminal
awplus(config)# interface vlan20
awplus(config-if)# ip igmp snooping mrouter interface port1.0.4
awplus(config-if)# ip igmp snooping mrouter interface port1.0.5
awplus(config-if)# ip igmp snooping mrouter interface port1.0.6
```

How Clients Leave Groups: Queries and Timers

When a client leaves a group, the Snoopers and the Querier check which ports now have clients that belong to that group. They will stop forwarding the group's traffic out any ports that are now unnecessary. In this section, we describe the process in detail.

Overview of leave process

The basic process when a client leaves a group is as follows:

1. The client sends a Leave/Done message to indicate that it no longer needs to receive that multicast group.
2. The Snooper receives the Leave/Done message and forwards it towards the Querier.
3. For all ports that belong to the group, the Querier changes its internal group membership timer to a short value (2 seconds by default—see "Querier timer values" below).
4. The Querier sends a Specific Query to ask which other clients still belong to that group.
5. The aforementioned Snooper receives the Specific Query. For all ports that belong to the group, the Snooper changes its internal group membership timer to a short value (2 seconds by default—see "Snooper timer values" below) unless the timer is already short. It forwards the Query out all its ports.
6. The Querier waits for the Last Member Query Interval time, 1 second by default, and then sends a second Specific Query.
7. The aforementioned Snooper snoops this second Specific Query and uses it to set the internal group membership timer for each port, unless the timer is already short (which it will be if the Snooper received the first Query). It forwards the Query out all its ports.
8. If the Snooper or Querier receives a Membership Report on a port, it sets the port timer to the "Group Membership Interval" value and continues to forward the multicast stream out that port. Otherwise, the timers for that port expire and the Snooper and/or Querier stops forwarding the multicast stream out that port.

Querier timer values

As described in Step 3 above, when the Querier sends a Specific Query for a group in response to a Leave message, the Querier updates a timer for ports that forward that group.

The timer is the following two values multiplied together:

- Last Member Query Count (LMQC)—the number of Specific Queries the Querier sends, 2 by default, and
- Last Member Query Interval (LMQI)—the time between the Specific Queries, 1 second by default

The default LMQC and LMQI give a timeout of 2 seconds. Therefore, by default the Querier must see the client response within 2 seconds of sending the first Specific Query.

Because of this process, sensible values for LMQC and LMQI are essential. In most networks, the defaults are appropriate and you should not change them. If you need to change them, see "Last Member Query Count and Last Member Query Interval" on page 74.

The commands **show ip igmp** and **show ipv6 MLD interface** displays the timer for the most recently updated port as the group's Refresh Time.

Snooper timer values

As described in Step 5 above in the "Overview of leave process" section, when the Snooper receives a Specific Query from the Querier, it may update a timer for ports that forward that group. The following flow chart describes the decision-making process for updating the timer.

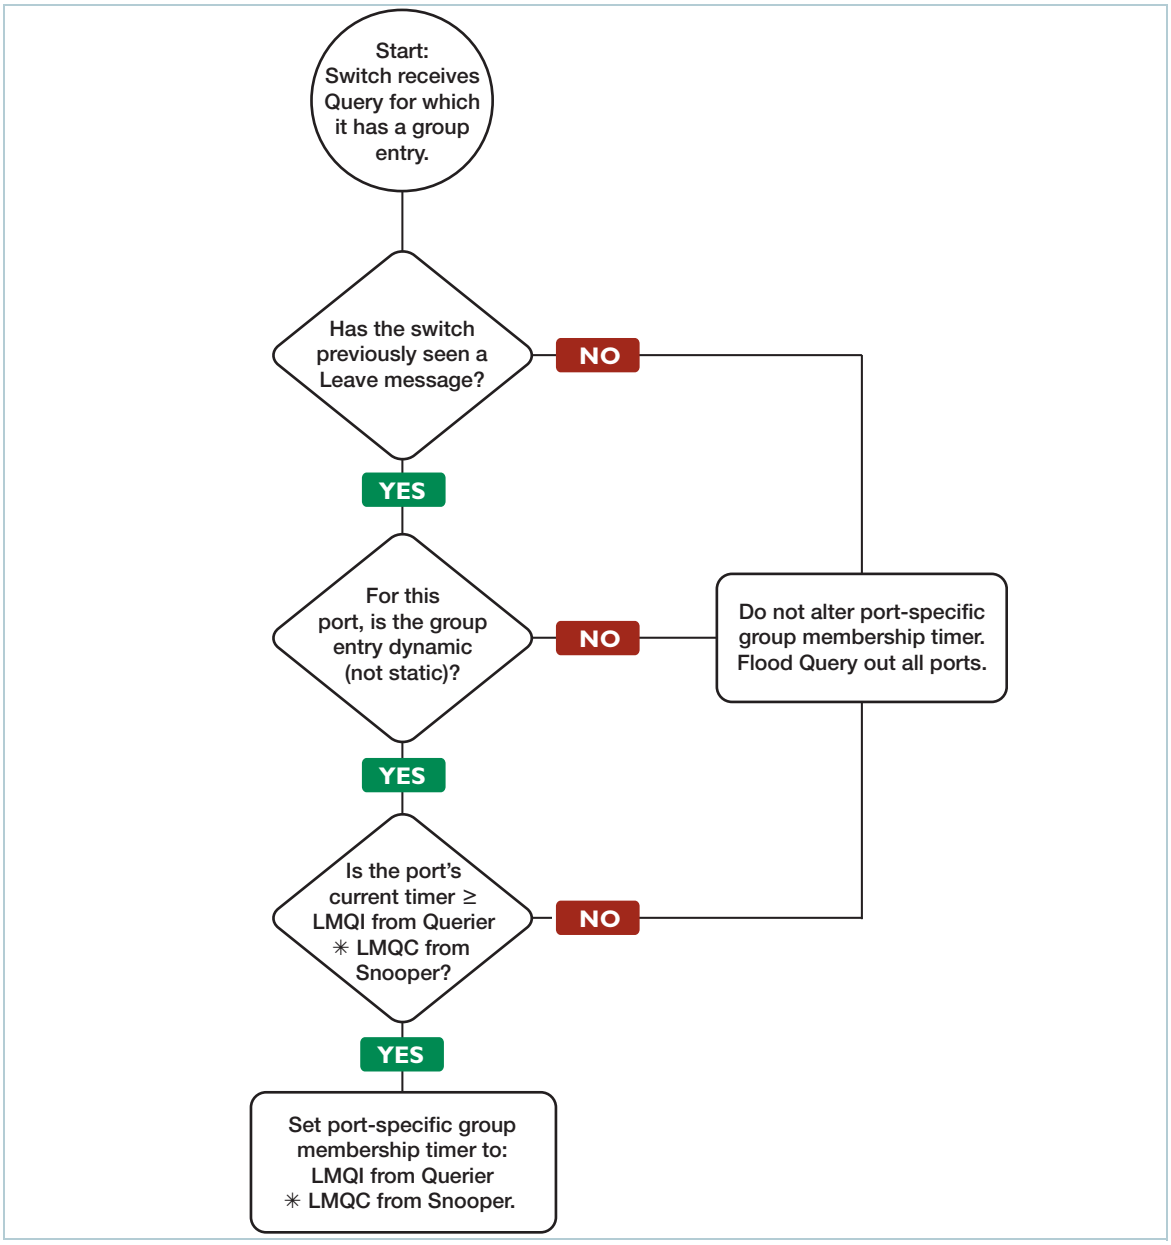

Note that:

- The command **show ip igmp snooping statistics interface** displays the timer for the most recently updated port as the time displayed in the entry:
  Group mode: Exclude (Expires: <***time***>).
- To calculate the timer, the Snooper takes the LMQI value that it receives from the Querier and multiplies it by the Snooper's own LMQC.
- The Snooper only reduces the timer if it receives a Leave message followed by a Specific Query—one of the messages is not enough.

Comparing the Querier and Snooper timers

By default, the Querier and Snooper port-specific group timers have the same value (2 seconds). This is because the LMQC is the same for the Querier and the Snoopers.

Consequences for high-loss and high-lag networks

If packet loss or lag time is an issue in your network, we recommend increasing the Robustness Variable on the Snoopers and the Querier.

On Allied Telesis Snoopers and Queriers, LMQC = Robustness Variable. For Snoopers, not all vendors make these counters the same. RFCs 2236 and 2710 recommend that LMQC and Robustness Variable have the same value on Queriers, but the IGMP/MLD timer rules for IGMP/MLD Snoopers are less well-defined.

Increasing the Querier LMQC (or Robustness Variable) increases the number of Specific Queries that the Querier sends. This increases the probability that an interested client will receive a Query.

Increasing the Snooper LMQC (or Robustness Variable) increases the length of time that the Snooper waits before aging out the port. This gives the client more time to reply to the Queries.

For example, if you increase the LMQC to 5 (the maximum) on the Querier and the Snooper, then the Querier sends 5 Queries and the Snooper waits for 5 * LMQI, which is 5 seconds with the default LMQI.

Make sure that the values on the Querier and the Snoopers match, so that the Snooper has time to forward all the Queries. For example, if you changed the Querier's Robustness Variable to 5 but left the Snooper unchanged, the Querier would send out 5 Queries 1 second apart but the Snooper would age out the group entry after only the first 2 Queries.

For more information about setting the Robustness Variable, and the consequences of this, see "Robustness Variable" on page 75.

IGMP/MLD Fast Leave

IGMP/MLD Fast Leave enhances your control over router bandwidth. Enabling Fast Leave tells IGMP/MLD snooping to stop the transmission of a group multicast stream to a port as soon as it receives a Leave/Done message on that port. No timeouts are observed.

Ordinarily, when IGMP/MLD snooping sees a Leave message, it waits for a Membership Query message before setting the entry timeout to 2 seconds. Fast Leave tells IGMP/MLD to drop the entry from the port as soon as the Leave/Done message is seen. For this reason, Fast Leave should only be configured on interfaces that have one client per port.

This example shows how to enable fast-leave processing on a VLAN.

```
awplus# configure terminal
awplus(config)# interface vlan1
awplus(config-if)# ip igmp snooping fast-leave
awplus(config-if)#ipv6 mld snooping fast-leave
```

Example

This example uses the same network configuration as the IGMP Snooping "Configuration example" on page 12. For convenience, the diagram is reproduced below.

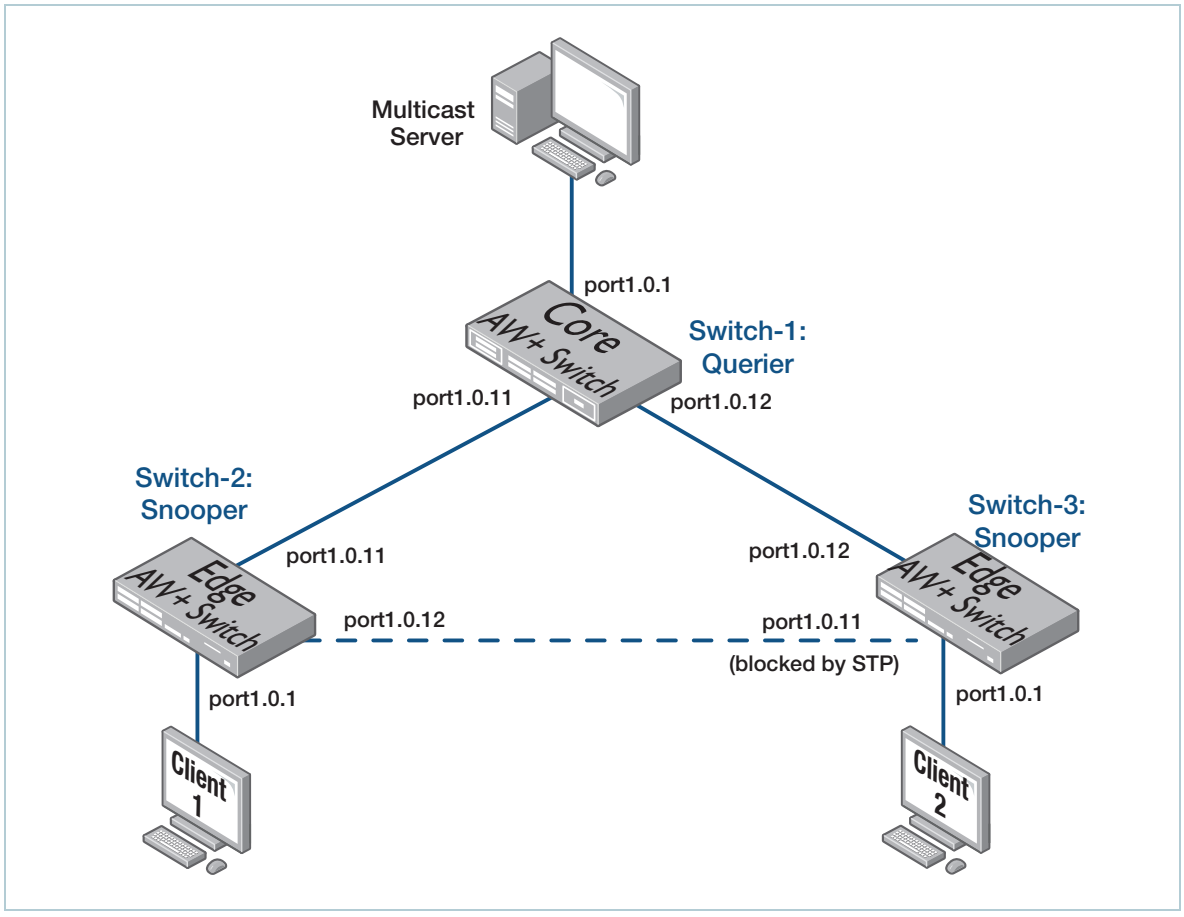

Step 1. **Configure Switch-1, the IGMP/MLD Querier**

```
!
hostname Switch-1
!
spanning-tree mode rstp
ipv6 forwarding
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
!
interface vlan100
 ip address 192.168.100.254/24
 ip igmp
 ip igmp version 2
 ipv6 address 2001:0db8::12:252/64
 ipv6 enable
 ipv6 mld
 ipv6 mld version 1
```

Step 2. **Configure Switch-2, an IGMP/MLD Snooper**

IGMP/MLD snooping is enabled by default and does not need any configuration.

```
!
hostname Switch-2
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
```

Step 3. **Configure Switch-3, also an IGMP/MLD Snooper**

Fast leave is enabled on this switch.

```
!
hostname Switch-3
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
!
interface vlan100
 ip igmp snooping fast-leave
 ipv6 mld snooping fast-leave
```

Explanation of IGMP/MLD Fast Leave

Imagine that Client-2 on Switch-3 sends a Membership Report to join the group 224.12.13.14 or FF02:10::1. The Snooper, Switch-3, adds this to its IGMP or MLD snooping table. When the same client then sends a Leave or Done message, the IGMP or MLD Querier responds with a Membership Query and waits for a configured time for a response.

The following sections describe in detail the differences between having Fast Leave disabled and enabled, but in summary:

- Without Fast Leave, the IGMP/MLD Snooper waits the same length of time as the Querier, then expires the entry if there was no response.
- With Fast Leave, the IGMP/MLD Snooper expires the entry as soon as it sees the Leave/Done message from the client. By the time the Querier sends the Membership Query, the Snooper will have already expired the entry and therefore stopped sending the stream to the client.

When Fast Leave is disabled

The IGMP/MLD Snooper sees the Membership Query from the Querier and accordingly sets its expiry time in seconds to match the Querier.

```
Switch-3# sh ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Disabled, Active(Snoop Only), Non-Querier, Version 3 (default)
 IGMP interface has 2 group-record states
 IGMP activity: 68 joins, 0 leaves
 IGMP querying router is 192.168.100.254
 IGMP robustness variable is 2
 IGMP last member query count is 2
 IGMP query interval is 125 seconds
 IGMP Startup query interval is 31 seconds
 IGMP Startup query count is 2
 IGMP query holdtime is 500 milliseconds
 IGMP querier timeout is 255 seconds
 IGMP max query response time is 10 seconds
 Group Membership interval is 260 seconds
 IGMP Last member query count is 2
 Last member query response interval is 1000 milliseconds
 Strict IGMPv3 ToS checking is disabled on this interface
 Source Address checking is enabled
 IGMP Snooping is globally enabled
 IGMP Snooping query solicitation is globally enabled for Root/Master Nodes
 IGMP Snooping query solicitation is globally disabled for Non Root/Master Nodes
  Num. query-solicit packets: 0 sent, 0 recvd
 IGMP Snooping is enabled on this interface
 IGMP Snooping fast-leave is not enabled
 IGMP Snooping querier is not enabled
 IGMP Snooping report suppression is enabled
```

If no Membership Report is received by the time the counters go to zero, then the client's entry is dropped from both the Querier and Snooper.

When you enable fast leave on Switch-3

When Fast Leave is enabled on Switch-3, but not on Switch-1, an interesting chain of events occurs when the client sends a Leave message, as shown in the following diagram.

The result of this is that Switch-3 adds the group back into its snooping table (with the same timeout as the IGMP Querier) but has no ports interested in receiving the group. Because Fast Leave is enabled on vlan100, that port was removed from the group as soon as Switch-3 received the Leave message.

```
Switch-3#sh ip igmp snooping statistics interface vlan100
IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         02:24:16
Group mode:     Include ()
Last reporter:  192.168.100.254

Port member list:
port1.0.12 - 135 secs

Interface:      vlan100
Group:          224.12.13.14
Flags:
Uptime:         00:12:23
Group mode:     Exclude (Expires: 00:02:19)
Last reporter:  192.168.100.20
Source list is empty

Port member list:
port1.0.1 - 140 secs
```

We can see the entry for 224.12.13.14 with port 1.0.1 associated with this entry. Now we send a Leave message from the client attached to port 1.0.1, and with Fast Leave enabled on vlan100 the switch now removes the port association while keeping the entry in the table as an unregistered group.

```
Switch-3#sh ip igmp snooping statistics interface vlan100
IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         02:24:30
Group mode:     Include ()
Last reporter:  192.168.100.254

Port member list:
port1.0.12 - 258 secs

Group Type:     UnReg MC Group
Interface:      vlan100
Group:          224.12.13.14
Uptime:         00:00:03
Group mode:     Include (Expires: 00:04:16)
Last reporter:  192.168.1.100
```

Adding Fast Leave to Switch-1 would not be sensible, since there may be clients attached to other ports on Switch-3. If you enabled Fast Leave on Switch-1, one Leave message from Switch-3 would drop the multicast stream for everyone on that switch.

Immediate Leave and Fast Leave

Both Layer 2 IGMP/MLD Immediate Leave and Layer 3 IGMP/MLD Snooping Fast Leave are enabled on a per interface basis.

The following example shows how to enable the Immediate Leave feature on the VLAN interface vlan2 for a specific range of multicast groups.

Configure IGMP/MLD Immediate Leave:

```
awplus# configure terminal
awplus(config)# access-list 34 permit 225.192.20.0 0.0.0.25
awplus(config)# ipv6 access-list standard FastGroup permit FF02:10::/48
awplus(config)# interface vlan2
awplus(config-if)# ip igmp immediate-leave group-list 34
awplus(config-if)# ipv6 mld immediate-leave group-list FastGroup
awplus(config-if)# exit
```

Configure IGMP/MLD Snooping Fast Leave:

```
awplus(config)# int vlan2
awplus(config-if)# ip igmp snooping fast-leave
awplus(config-if)# ipv6 mld snooping fast-leave
awplus(config-if)# exit
awplus(config)# sh run int vlan2
```

```
!
interface vlan2
 ip igmp snooping fast-leave
 ipv6 mld snooping fast-leave
!
```

Configurable IGMP/MLD Timers and Counters

This section looks at some of the timers and counters that control how often IGMP/MLD sends queries and how quickly entries time out. First, it gives background information in the following subsections:

Then it looks at each of the configurable timers and counters, in the following subsections:

RFC 2236 also describes other counters and timers that this section does not describe, because this section only describes the counters and timers that you can directly set. The switch derives other counters and timers from the above subset.

Timer and counter relationships

The above timers and counters are related to each other and to others in RFCs 2236 and 270 by the following formulae:

- Last Member Query Count = Robustness Variable (note, in RFC 2710, the name Last Listener Query Count is used, rather than Last Member Query Count).
- Group Membership Interval = (Robustness Variable * Default Query Interval) + one Query Response Interval in seconds
- Startup Query Count = Robustness Variable
  The Startup Query Count is the number of General Queries that the Querier sends when it starts up.
- Other Querier Timeout = (Robustness Variable * Default Query Interval) + (Query Response Interval in seconds/2)
  The Other Querier Timeout is the length of time that a potential Querier waits after receiving a Query before it assumes that it should become the Querier.

These relationships mean you need to take care if you change timers or counters. "Stopping Snoopers from Snooping Non-IGMP Messages" on page 82 describes the consequences of a bad combination of values.

Default values

To display the values of the configurable IGMP settings for a VLAN, use the command:

```
awplus# show ip igmp int vlan<vid>
```

The following output shows the default values:

```
Switch-1#show ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Enabled, Active, Querier, Configured for version 2
 Internet address is 192.168.100.254
 IGMP interface has 2 group-record states
 IGMP activity: 1932 joins, 8 leaves
 IGMP robustness variable is 2
 IGMP last member query count is 2
 IGMP query interval is 125 seconds
 IGMP Startup query interval is 31 seconds
 IGMP Startup query count is 2
 IGMP query holdtime is 500 milliseconds
 IGMP querier timeout is 255 seconds
 IGMP max query response time is 10 seconds
 Last member query response interval is 1000 milliseconds
 Group Membership interval is 260 seconds
 IGMP Last member query count is 2
 Strict IGMPv3 ToS checking is disabled on this interface
 Source Address checking is enabled
 IGMP Snooping is globally enabled
 IGMP Snooping query solicitation is globally enabled for Root/Master Nodes
 IGMP Snooping query solicitation is globally disabled for Non Root/Master Nodes
  Num. query-solicit packets: 0 sent, 0 recvd
 IGMP Snooping is enabled on this interface
 IGMP Snooping fast-leave is not enabled
 IGMP Snooping querier is not enabled
 IGMP Snooping report suppression is enabled
```

Similarly, **show ipv6 MLD interface** outputs the default values for MLD:

```
awplus#show ipv6 mld interface
Interface vlan1 (Index 301)
MLD Enabled, Active, Querier, Version 2 (default)
Internet address is fe80::215:77ff:fec9:7468
MLD interface has 0 group-record states
MLD activity: 0 joins, 0 leaves
MLD robustness variable is 2
MLD last member query count is 2
MLD query interval is 125 seconds
MLD querier timeout is 255 seconds
MLD max query response time is 10 seconds
Last member query response interval is 1000 milliseconds
Group Membership interval is 260 seconds
MLD Snooping is globally enabled
MLD Snooping is enabled on this interface
MLD Snooping fast-leave is not enabled
MLD Snooping querier is enabled
MLD Snooping report suppression is enabled
```

All changes are made on the VLAN interface in which the user wants to make the changes:

```
Switch-1(config)#int vlan100
Switch-1(config-if)#ip igmp ?
  access-group                IGMP group access group
  immediate-leave             Leave groups immediately without
                              sending last member query, use for one
                              host  network only
  last-member-query-count     Last Member Query Count
  last-member-query-interval  Last Member Query Interval
  limit                       IGMP Limit
  mroute-proxy                Mroute to IGMP proxy
  proxy-service               Enable IGMP mroute proxy service
  querier-timeout             IGMP previous querier timeout
  query-holdtime              Query Hold Time
  query-interval              Query Interval
  query-max-response-time     IGMP Max Query Response Time
  ra-option                   IGMP Strict RA Option Validation
  robustness-variable         Robustness Variable
  snooping                    Layer 2 Snooping
  source-address-check        Enforce report source address checking
  startup-query-count         Startup Query count
  startup-query-interval      Startup Query Interval
  static-group                Static Group to be Joined
  v3toscheck                  IGMPv3 strict IP ToS checking (0xC0)
  version                     IGMP Version
```

Similarly, MLD supports commands to alter the following:

- last-member-query-count
- last-member-query-interval
- querier-timeout
- query-interval
- query-max-response-time
- robustness-variable

However, it does not provide commands to alter startup-query-count or startup-query-interval. Note that units for Last Member Query Interval (LMQI) and Query Response Interval are .001 seconds (i.e., milliseconds). Therefore, the default LMQI of 1 second is shown as 1000ms.

The units for the Max Query Response Interval are seconds. The default Max Query Response Interval is 10 seconds.

Last Member Query Count and Last Member Query Interval

The Last Member Query Count (LMQC) is the number of Specific Queries the Querier sends after receiving a Leave message—1 to 5 messages.

The Last Member Query Interval (LMQI) is the time between the Specific Queries—1000 to 25500 in units of 0.001 seconds.

These counters determine how quickly a group times out when the last client leaves the group. You should read this section in conjunction with "How Clients Leave Groups: Queries and Timers" on page 61, which has an outline of the leave process and a detailed discussion of the timers.

What these counters do

On the Querier and all Snoopers, IGMP/MLD keeps group membership timeout values on each port. During general multicasting, these timeouts are (by default) 260 seconds. When a client leaves a group, these timeouts are reduced to make multicasting stop quickly after the last client leaves. The LMQC and LMQI determine the value of the timeout during this leave process (2 seconds with the default LMQC and LMQI).

- On the Querier:
  timeout during the leave process = LMQC * LMQI
- On Snoopers:
  timeout during the leave process = LMQI from Querier * LMQC from Snooper

The commands **show ip igmp interface vlan** and **show ipv6 mld interface** display the timers for a particular interface.

Potential problems with changing these counters

For most networks, the default LMQI and LMQC values work. You should only change them if you are aware of the likely effect on the network. In particular, note that:

- Changing the LMQC automatically changes the Robustness Variable. Therefore, we do not recommend setting the LMQC to 1, because it removes the system's allowance for packet loss. See "Robustness Variable" on page 75 for more information about the consequences of changing the Robustness Variable.
- If you set the LMQI (or LMQC, or both) too low, clients will not be able to reply to Specific Queries quickly enough and the Querier and Snoopers may delete group entries for ports that still need to receive multicasts. If this happens, some or all clients briefly lose the multicast stream.
- The default values of LMQI (1000ms) and LMQC (2) mean that the Querier must receive client Membership Reports within 2 seconds of the first Query. This is already quite a short time and we do not recommend reducing it even more. For example, reducing the LMQI to 500ms would allow only one second for responses, which may be too little.

How to change these counters

The following example increases LMQI a lot, and then shows the resulting changed refresh time.

```
Switch-1#configure terminal
Switch-1(config)#int vlan100
Switch-1(config-if)#ip igmp last-member-query-interval ?
 <1000-25500> Last Member Query Interval value (Default: 1000 ms)
Switch-1(config-if)#ip igmp last-member-query-interval 255
```

```
Switch-1#show ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Enabled, Active, Querier, Configured for version 2
 Internet address is 192.168.100.254
 IGMP interface has 2 group-record states
 IGMP activity: 1966 joins, 8 leaves
 IGMP robustness variable is 2
 IGMP last member query count is 2
 IGMP query interval is 125 seconds
 IGMP Startup query interval is 31 seconds
 IGMP Startup query count is 2
 IGMP query holdtime is 500 milliseconds
 IGMP querier timeout is 255 seconds
 IGMP max query response time is 10 seconds
 Last member query response interval is 255 milliseconds
 Group Membership interval is 260 seconds
 IGMP Last member query count is 2
 Strict IGMPv3 ToS checking is disabled on this interface
 Source Address checking is enabled
 IGMP Snooping is globally enabled
 IGMP Snooping query solicitation is globally enabled for Root/Master Nodes
 IGMP Snooping query solicitation is globally disabled for Non Root/Master Nodes
  Num. query-solicit packets: 0 sent, 0 recvd
 IGMP Snooping is enabled on this interface
 IGMP Snooping fast-leave is not enabled
 IGMP Snooping querier is not enabled
 IGMP Snooping report suppression is enabled
```

The refresh time in seconds is (LMQI/1000) * LMQC = 255/1000 * 2 = 0.5 seconds

Robustness Variable

What this counter does

The Robustness Variable (RV) allows you to tune for the expected packet loss on a subnet. If you expect a subnet to be lossy, you can increase the RV. IGMP/MLD is robust to packet loss of one packet less than the RV. The RV is an integer from 1 to 5 and should not be set to 1. If packet loss or lag time is an issue in your network, we recommend increasing the Robustness Variable on the Snoopers and the Querier. This increases the:

- number of Queries that the Querier sends out (by increasing the LMQC)
- amount of time that the Querier and the Snoopers wait for clients to reply

For more details, see "Consequences for high-loss and high-lag networks" on page 64.

Potential problems with changing this counter

The RV is the counter you are most likely to need to change. However, you need to appreciate the effect this has on your network, as described in RFC 2236 and RFC 2710 Changing the RV changes other values from the RFC, as follows:

- Last Member Query Count = Robustness Variable
- Group Membership Interval = (Robustness Variable * Default Query Interval) + one Query Response Interval in seconds
- Startup Query Count = Robustness Variable
- Other Querier Timeout = (Robustness Variable * Default Query Interval) + (Query Response Interval in seconds/2)

The **Group Membership Interval** automatically changes to match the above formula.

- Group Membership Interval = (Robustness Variable * Default Query Interval) + one Query Response Interval in seconds

How to change this counter

```
Switch-1# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch-1(config)# int vlan100
Switch-1(config-if)# ip igmp robustness-variable ?
 <1-7> Robustness Variable value (Default: 2)
Switch-1(config-if)# ip igmp robustness-variable 5
```

Note that the Last Member Query Count, Startup Query Count, Group Membership Interval and Last Member Query Count have also changed.

```
Switch-1#sh ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Enabled, Active, Querier, Configured for version 2
 Internet address is 192.168.100.254
 IGMP interface has 3 group-record states
 IGMP activity: 18 joins, 0 leaves
 IGMP robustness variable is 5
 IGMP last member query count is 5
 IGMP query interval is 125 seconds
 IGMP Startup query interval is 31 seconds
 IGMP Startup query count is 5
 IGMP query holdtime is 500 milliseconds
 IGMP querier timeout is 630 seconds
 IGMP max query response time is 10 seconds
 Last member query response interval is 1000 milliseconds

output is continued on next page...
```

```
Group Membership interval is 635 seconds
 IGMP Last member query count is 5
 Strict IGMPv3 ToS checking is disabled on this interface
 Source Address checking is enabled
 IGMP Snooping is globally enabled
 IGMP Snooping query solicitation is globally enabled for Root/Master Nodes
 IGMP Snooping query solicitation is globally disabled for Non Root/Master Nodes
  Num. query-solicit packets: 0 sent, 0 recvd
 IGMP Snooping is enabled on this interface
 IGMP Snooping fast-leave is not enabled
 IGMP Snooping querier is not enabled
 IGMP Snooping report suppression is enabled
```

Default Query Interval

What this timer does

To maintain an accurate picture of group membership, the Querier periodically sends General Queries to all its IGMP interfaces. The Default Query Interval is the gap between General Queries.

Note that General Queries are quite different from Specific Queries, which the Querier sends to a group address when it receives a Leave message for that group.

The router only sends General or Specific Queries if it is the Querier. If another router is elected as the Querier, the non-elected router ignores the Default Query Interval and other such settings.

Potential problems with changing this timer

If you change the Default Query Interval, the Group Membership Interval also needs to change so that clients have an appropriate amount of time to reply to the Query.

How to change this timer

The Default Query Interval is an integer from 1 to 65535 seconds, specified using the **query-interval** parameter. The default is 125 seconds. The following example tweaks the interval. Note that the Querier timeout and Group Membership Interval also change.

```
Switch-1#configure terminal
Switch-1(config)#int vlan100
Switch-1(config-if)#ip igmp quer
Switch-1(config-if)#ip igmp query-interval?
 <2-18000> Query Interval value (Default: 125 s)
Switch-1(config-if)#ip igmp query-interval 120
```

```
Switch-1#sh ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Enabled, Active, Querier, Configured for version 2
 Internet address is 192.168.100.254
 IGMP interface has 3 group-record states
 IGMP activity: 16 joins, 0 leaves
 IGMP robustness variable is 2
 IGMP last member query count is 2
 IGMP query interval is 120 seconds
 IGMP Startup query interval is 31 seconds
 IGMP Startup query count is 2
 IGMP query holdtime is 500 milliseconds
 IGMP querier timeout is 245 seconds
 IGMP max query response time is 10 seconds
 Last member query response interval is 1000 milliseconds
 Group Membership interval is 250 seconds
 IGMP Last member query count is 2
 Strict IGMPv3 ToS checking is disabled on this interface
 Source Address checking is enabled
 IGMP Snooping is globally enabled
 IGMP Snooping query solicitation is globally enabled for Root/Master Nodes
 IGMP Snooping query solicitation is globally disabled for Non Root/Master Nodes
  Num. query-solicit packets: 0 sent, 0 recvd
 IGMP Snooping is enabled on this interface
 IGMP Snooping fast-leave is not enabled
 IGMP Snooping querier is not enabled
 IGMP Snooping report suppression is enabled
```

Max Query Response Interval

What this timer does

The Max Query Response Interval determines the longest time clients can take to reply to a General Query. The Querier inserts the Max Query Response Interval into General Query messages. Clients randomly choose a time between 0 and the Max Query Response Interval at which to respond to a General Query. Increasing the Query Response Interval spreads IGMP messages over a longer time period, which reduces the burstiness of traffic on the network.

It may be useful to decrease the Max Query Response Interval if you are running EPSR or RSTP. Decreasing the Max Query Response Interval reduces the recovery time after a topology change. For more information, see "Query Solicitation - Rapid Recovery From Topology Changes" on page 35.

Potential problems with changing this timer

If your network has many multicast clients and you make the Max Query Response Interval too short, you may congest the Snoopers and Querier with too many Report messages in a short time.

If you change the Max Query Response Interval, the **Group Membership Interval** also needs to change, so that clients have an appropriate amount of time to reply to the Query.

How to change this timer

The Max Query Response Interval is an integer from 1 to 3180 seconds, specified using the **query-max-response-time** parameter. The default is 10 seconds.

```
Switch-1# configure  terminal
Switch-1(config)# int vlan100
Switch-1(config-if)# ip igmp query-max-response-time ?
 <1-3180> Query Response Time (Default: 10 s)
Switch-1(config-if)#ip igmp query-max-response-time 100
```

```
Switch-1#sh ip igmp int vlan100

Interface vlan100 (Index 400)
 IGMP Enabled, Active, Querier, Configured for version 2
 Internet address is 192.168.100.254
 IGMP interface has 3 group-record states
 IGMP activity: 12 joins, 0 leaves
 IGMP robustness variable is 2
 IGMP last member query count is 2
 IGMP query interval is 125 seconds
 IGMP Startup query interval is 31 seconds
 IGMP Startup query count is 2
 IGMP query holdtime is 500 milliseconds
 IGMP querier timeout is 300 seconds
 IGMP max query response time is 100 seconds
 Last member query response interval is 1000 milliseconds
 Group Membership interval is 350 seconds
 IGMP Last member query count is 2
 Strict IGMPv3 ToS checking is disabled on this interface
 Source Address checking is enabled
 IGMP Snooping is globally enabled
 IGMP Snooping query solicitation is globally enabled for Root/Master Nodes
 IGMP Snooping query solicitation is globally disabled for Non Root/Master Nodes
  Num. query-solicit packets: 0 sent, 0 recvd
 IGMP Snooping is enabled on this interface
 IGMP Snooping fast-leave is not enabled
 IGMP Snooping querier is not enabled
 IGMP Snooping report suppression is enabled
```

In this example, the Querier sends a General Query every 125 seconds. The General Query contains the Max Query Response Interval of 100, which tells clients that they have 100 seconds to reply to this General Query message.

Group Membership Interval

What this timer does

The Group Membership Interval specifies the length of time before the switch deletes a group from its multicast group database after the switch last receives a Membership Report for that group. All IGMP/MLD routers and switches in a network use this interval to maintain their group membership databases, not just the Querier.

The Querier also uses this interval to close down multicasting if it receives no replies to all General Queries for a group. If the Querier sends a General Query and does not receive any Membership Reports in response, it continues to send any existing multicast streams. In the meanwhile, the

Group Membership Interval counts down until it hits zero, at which point the Querier stops propagating the multicast streams through the LAN.

Potential problems with changing this timer

Making the Group Membership Interval too short has serious consequences. You remove the network's ability to cope with losing a General Query and you may not give enough time for client responses to reach the Querier. These problems would cause multicasting to stop for some or all clients. For more information, see "Stopping Snoopers from Snooping Non-IGMP Messages" on page 82.

Synchronization of timers

The Group Membership Interval is a function of several other timers, according to the following formula from RFC 2236:

Group Membership Interval = (Robustness Variable * Default Query Interval) + one Query Response Interval in seconds

The Group Membership Interval changes automatically if you change either the Robustness Variable or the Query Response Interval.

Table 5: Summary of the Timers and Counters used in MLDv1

PARAMETER	MEANING	DEFAULT VALUE	RELATIONSHIP TO OTHER PARAMETERS	ALLIEDWARE PLUS CONFIGURATION COMMAND
Query Interval	Period at which querier sends General Queries	125 sec	Must be greater than the Query Response Interval	IPv6 MLD query interval (per interface)
Query Response Interval	The value put into the Maximum Response Delay of General Query packets	10 sec	Must be less than the Query Interval	Not configurable
Multicast Listener Interval	The period, during which no reports are received for a given Group on a given interface, before the querier deletes the forwarding entry for that Group on that interface	Defined by relationship to other parameters	Defined as Robustness Variable X Query Interval + Query Response Interval	Not configurable
Other Querier Present Interval	The period, during which no Queries, with a lower source IPv6 address, are received on a given interface, before a device transitions from nonquerier to Active Querier on that interface.	The RFC states that this is defined by its relationship to other parameters. AlliedWare Plus sets the default to 255 seconds	Should be Robustness Variable X Query Interval + (Query Response Interval / 2)	IPv6 MLD querier timeout (per interface)

Table 5: Summary of the Timers and Counters used in MLDv1 (continued)

PARAMETER	MEANING	DEFAULT VALUE	RELATIONSHIP TO OTHER PARAMETERS	ALLIEDWARE PLUS CONFIGURATION COMMAND
Startup Query Interval	The time between the multiple General Queries sent on a newly active Querier interface.	Query Interval /4	Should be Query Interval / 4	Not configurable
Startup Query Count	Number of General Queries sent when a querier interface starts up	Equal to Robustness Variable	Should be the same as Robustness Variable	Not configurable
Last Listener Query Count	The number of Multicast-Address Specific Queries sent after a Listener Done message is received	Equal to Robustness Variable	Should be the same as Robustness Variable	IPv6 MLD last member query count (per interface)
Last Listener Query Interval	The time between each of the Multicast Address-Specific Queries sent after a Listener Done message is received.	1000	N/A	IPv6 MLD last member query count (per interface)
Robustness Variable	The number of times that certain packets are repeated, to allow for the possibility that packets will be lost.	2	Used to calculate the expected values of several other parameters, as seen above.	ipv6 mld robustness variable (per interface)
Unsolicited Report Interval	When a host first wishes to receive a given group, it sends out a number (equal to the Robustness Variable) of unsolicited reports to request this group. This parameter is the time between these reports.	10 seconds	N/A	N/A

Stopping Snoopers from Snooping Non-IGMP Messages

Some networks include routers that have no interest in IGMP, but still generate multicast messages by running protocols like OSPF. When a Snooper receives multicast messages from such a router, the Snooper adds the router's port to its All Groups port list. This means the router is unnecessarily sent IGMP and multicast traffic. Using IGMP features to prevent the excess traffic is particularly helpful when you cannot or do not want to control the traffic at the router.

This example describes how to use AlliedWare Plus' advanced IGMP features to prevent this behavior, by limiting the ports that the Snooper adds to the All Groups list, or by stopping particular types of traffic from adding ports to the All Groups list.

Example

The example is based around a network that has a router running OSPF. The router is connected to a LAN through a switch. The LAN is a single subnet with no OSPF routers inside it. The network for this example uses the same loop as for "IGMP/MLD Snooping" on page 10, with a router attached to Switch-2. The network is shown in the following figure:

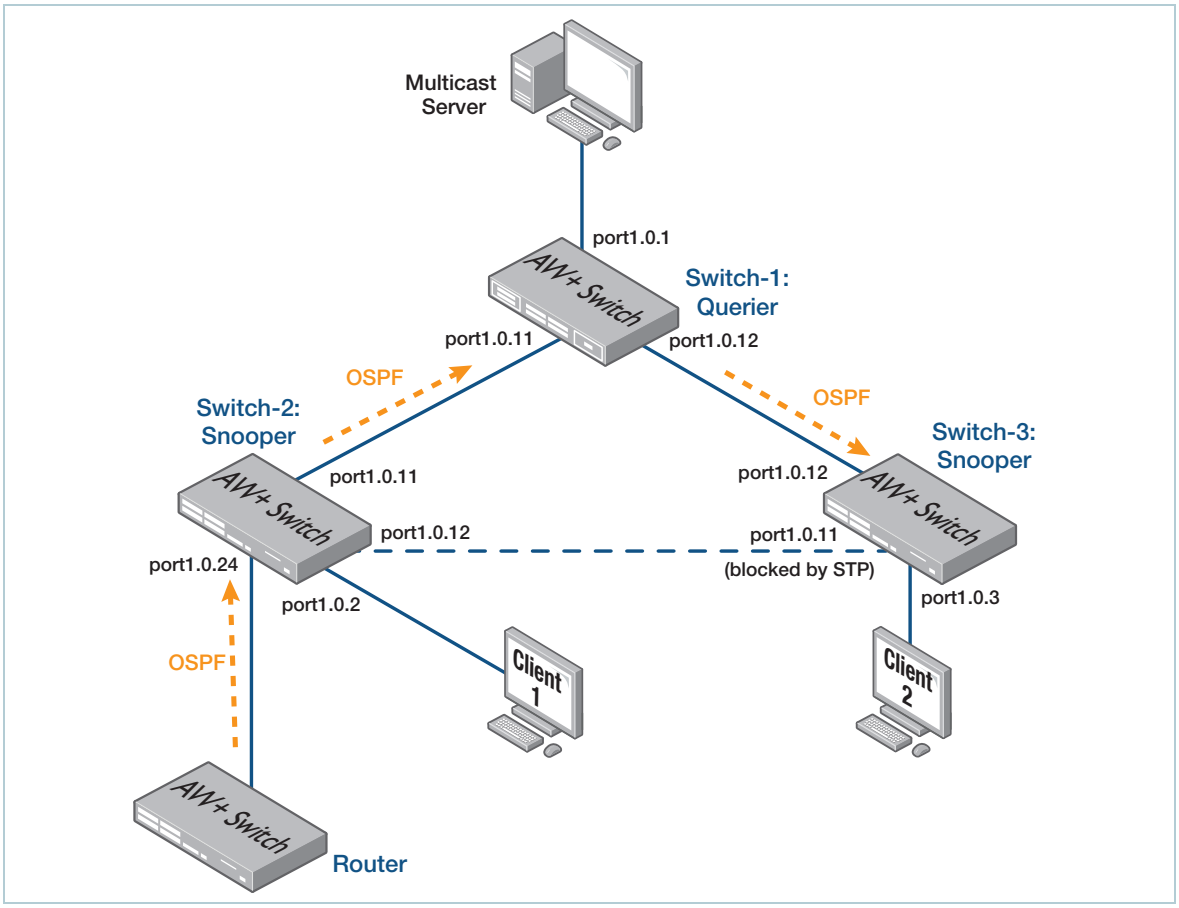

Each example in this section modifies the following base configuration.

Switch-1 Configuration—IGMP Querier

```
!
hostname Switch-1
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
 vlan 100 state enable
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
!
interface vlan100
 ip address 192.168.100.254/24
 ip igmp
 ip igmp version 2
```

Switch-2 Configuration—IGMP Snooper

```
!
hostname Switch-2
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
 vlan 100 state enable
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
```

Switch-3 Configuration—IGMP Snooper

```
!
hostname Switch-3
!
spanning-tree mode rstp
!
vlan database
 vlan 100 name vlan100
 vlan 100 state enable
!
interface port1.0.1-1.0.10
 switchport access vlan 100
 spanning-tree edgeport
!
interface port1.0.11-1.0.12
 switchport
 switchport mode access
 switchport access vlan 100
!
interface port1.0.13-1.0.24
 switchport access vlan 100
 spanning-tree edgeport
```

Router Configuration—OSPF

```
set system name=Router

# VLAN general configuration
create vlan=vlan100 vid=100
add vlan=100 port=1-4
create vlan=vlan2 vid=2
add vlan=2 port10-12

# IP configuration
enable ip
set ip autonomous=65000
add ip int=vlan100 ip=192.168.100.100 mask=255.255.255.0
add ip int=vlan2 ip=10.0.0.1

# OSPF configuration
set ospf routerid=192.168.100.100
add ospf area=0.0.0.1
add ospf range=192.168.100.0 area=0.0.0.1 mask=255.255.255.0
add ospf interface=vlan100 area=0.0.0.1
enable ospf
```

With the configuration above, the router sends OSPF messages to Switch-2. As the following outputs show, this means that:

- Switch-2 designates port 1.0.11 and 1.0.24 as the Router ports.
- Switch-2 forwards the OSPF packets to port 1.0.11 on Switch-1, so Switch-1 designates port1.0.11 as the Router port.
- Switch-1 forwards the OSPF packets to uplink port 1.0.12 on Switch-3, so Switch-3 designates port1.0.12 as a Router port.

```
Switch-1# sh ip igmp snooping statistics interface vlan100
IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         00:07:07
Group mode:     Include ()
Last reporter:  192.168.100.100

Port member list:
port1.0.11 - 250 secs
```

Switch-2 has two entries for the Router Port: port1.0.11, which is connected to Switch-1—the IGMP Querier—and port1.0.24, which has the OSPF router connected to it.

```
Switch-2#sh ip igmp snooping statistics interface vlan100
IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         00:16:09
Group mode:     Include ()
Last reporter:  192.168.100.100

Port member list:
port1.0.24 - 251 secs
port1.0.11 - 215 secs
```

On Switch-3, the Router port is learnt on interface port1.0.12

```
Switch-3#sh ip igmp snooping statistics interface vlan100
IGMP Snooping statistics for vlan100
Group Type:     Router Port Learnt
Interface:      vlan100
Group:          224.0.0.2
Uptime:         00:08:32
Group mode:     Include ()
Last reporter:  192.168.100.100

Port member list:
port1.0.12 - 250 secs
```

Controlling which addresses create All Groups entries

The router or switch adds a port to its All Groups list when it determines that the port has a router attached to it. This example shows how to influence the switch's process in determining who is a router, and therefore when to add a port to the All Groups list.

You can control the criteria for deciding which packets actually indicate the presence of a router on a particular port, by using the command:

```
awplus(config-if)# ip igmp snooping routermode {all|default|ip|
multicastrouter|address <ip-address>}
```

With this command, you specify (in effect) a list of IP addresses. When the switch receives a multicast packet on a port, it compares the packet's destination IP address with the list. If they match, the switch considers the packet to be from a ‘router’, and adds the port to the All Groups list.

The following table shows the address lists that each command option gives.

THIS OPTION...	MEANS THAT THE PORT IS TREATED AS A MULTICAST ROUTER PORT IF IT RECEIVES PACKETS FROM...
all	any reserved multicast addresses (224.0.0.1 to 224.0.0.255)
default	224.0.0.1 (IGMP Queries) 224.0.0.2 (all routers on this subnet) 224.0.0.4 (DVMRP routers) 224.0.0.5 (all OSPF routers) 224.0.0.6 (OSPF designated routers) 224.0.0.9 (RIP2 routers) 224.0.0.13 (all PIM routers) 224.0.0.15 (all CBT routers)
multicastrouter	224.0.0.4 (DVMRP routers) 224.0.0.13 (all PIM routers)
ip	the current list of addresses, plus addresses specified using the command **ip igmp snooping routermode address**, minus addresses specified using the command **no ip igmp snooping routermode address**.

```
Switch-2# sh ip igmp snooping routermode
Router mode.............IP
Reserved multicast address
```

IGMP Packet Reception Control

AlliedWare Plus provides the ability to control which IGMP packets will be accepted on a given port. The command that implements this control is:

```
[no] ip igmp trusted [all|query|report|routermode]
```

where:

all allows IGMP to accept and act on all IGMP and other routermode packets (OSPF, RIP, PIM, etc)

query allows IGMP to accept and act on IGMP queries

report allows IGMP to accept and act on IGMP membership reports

routermode allows IGMP to accept and act on 'routermode' packets (OSPF, RIP, PIM, etc.)

This command is used for disallowing specified packets from being processed by the IGMP module if the packets are received on the specified ports/aggregator.

By default, all ports and aggregators are all trusted interfaces, ie. IGMP is allowed to process all IGMP query, report, and router mode packets arriving on all interfaces.

Example To disallow processing of IGMP query packets arriving on port1.0.5 by IGMP module, use the command:

```
awplus(config)#int port1.0.5
awplus((config-if)#no ip igmp trusted query
```

This is useful in cases where 'rogue queriers' exist on a network. If a snooping switch receives queries from rogue queriers, it will make the receiving ports router ports. So potentially, multicast would be flooded out a number of ports that it does not need to be. Configuring the snooper to ignore queries on certain ports will negate the effect of the rogue queriers.

IGMP Debugging

In this section, we shall examine the debugging messages that the switch outputs when certain events occur while debugging is enabled. To enable debugging, use the command:

```
awplus# debug igmp {all|decode|encode|events|fsm|tib}
```

```
Switch-1# debug igmp ?
 all     Turn on all Debugging
 decode  IGMP decode
 encode  IGMP encode
 events  IGMP events
 fsm     IGMP FSM
 tib     IGMP Tree-Info-Base (TIB)
```

PARAMETER	MEANING
all	Enable or disable all debug options for IGMP
decode	Debug of IGMP packets that have been received
encode	Debug of IGMP packets that have been sent
events	Debug IGMP events
fsm	Debug IGMP Finite State Machine (FSM)
tib	Debug IGMP Tree Information Base (TIB)

A client joins a group

Client 2 sends a Membership Report for group 224.12.13.14. Switch-1 sees the report on vlan100 (1.0.11), and adds the port to its IGMP and IGMP snooping tables.

In this example we will show this using the following debugging options:

```
debug igmp decode
debug igmp events
```

```
Switch-1#19:52:20 Switch-1 IMISH[26241]: debug igmp decode
19:52:25 Switch-1 NSM[1257]: [IGMP-DECODE] : IP packet 192.168.100.100->224.0.0.5 on
vlan100(port1.0.11) triggers router port detection
19:52:25 Switch-1 NSM[1257]: [IGMP-DECODE] : UnReg Grp 224.12.13.14 on vlan100
19:52:34 Switch-1 NSM[1257]: [IGMP-DECODE] : IGMP V2 Membership Report, (Src=0.0.0.0) Max. Rsp.
Code 0 on port1.0.11
19:52:34 Switch-1 NSM[1257]: [IGMP-DECODE] : Grp 224.12.13.14 on vlan100
```

```
Switch-1#term mon
Switch-1#debug igmp events
```

```
19:59:05 Switch-1 IMISH[2155]: debug igmp events
19:59:15 Switch-1 NSM[1260]: [IGMP-EVENTS] : Deleting Unreg MC Grp 224.12.13.14 on vlan100
19:59:15 Switch-1 NSM[1260]: [IGMP-EVENTS] : Grp Timer Refresh 224.12.13.14 on
vlan100(port1.0.11)
19:59:16 Switch-1 NSM[1260]: [IGMP-EVENTS] : Expiry (Host-Side Generate Grp Report Timer)  for
224.12.13.14 on vlan100
19:59:17 Switch-1 NSM[1260]: [IGMP-EVENTS] : Expiry (Host-Side Generate Grp Report Timer)
for 224.12.13.14 on vlan100
19:59:21 Switch-1 NSM[1260]: [IGMP-EVENTS] : Exipry on vlan100
19:59:22 Switch-1 NSM[1260]: [IGMP-EVENTS] : Grp Timer Refresh 224.12.13.14 on
vlan100(port1.0.11)
```

A client leaves a group

Client 2 sends a Leave message for group 224.12.13.14. Switch-1 sees the Leave message on vlan100 (1.0.11).

```
Switch-1#19:52:20 Switch-1 IMISH[26241]: debug igmp decode
20:04:25 Switch-1 NSM[1260]: [IGMP-DECODE] : IGMP V2 Membership Report, (Src=0.0.0.0) Max. Rsp.
Code 0 on port1.0.11
20:04:25 Switch-1 NSM[1260]: [IGMP-DECODE] : Grp 224.12.13.14 on vlan100
20:04:25 Switch-1 NSM[1260]: [IGMP-DECODE] : IP packet 192.168.100.100->224.0.0.5 on
vlan100(port1.0.11) triggers router port detection
20:04:33 Switch-1 NSM[1260]: [IGMP-DECODE] : IGMP V2 Leave Group, (Src=0.0.0.0) Max. Rsp.
Code 0 on port1.0.11
20:04:33 Switch-1 NSM[1260]: [IGMP-DECODE] : Grp 224.12.13.14 on vlan100
20:04:35 Switch-1 NSM[1260]: [IGMP-DECODE] : UnReg Grp 224.12.13.14 on vlan100
```

We can see the Membership Report for group 224.12.13.14 on port1.0.11, without a source IP address.

And then we can see the Leave received on port 1.0.11.

```
Switch-1# term mon
Switch-1# debug igmp events
```

```
Switch-1#20:02:32 Switch-1 NSM[1260]: [IGMP-EVENTS] : Grp Timer Refresh 224.12.13.14 on
vlan100(port1.0.11)
20:02:44 Switch-1 NSM[1260]: [IGMP-EVENTS] : Grp Timer Refresh 224.12.13.14 on
vlan100(port1.0.11)
20:02:45 Switch-1 NSM[1260]: [IGMP-EVENTS] : Expiry (Group Query Re-Transmit Timer) for Grp
224.12.13.14 on vlan100
20:02:46 Switch-1 NSM[1260]: [IGMP-EVENTS] : Expiry (Group Liveness Timer) for Grp 224.12.13.14
on vlan100
20:02:46 Switch-1 NSM[1260]: [IGMP-EVENTS] : Expiry (L2 IF Grp Liveness Timer)  for Grp
224.12.13.14 on vlan100(port1.0.11)
20:02:46 Switch-1 NSM[1260]: [IGMP-EVENTS] : Setting Grp 224.12.13.14 on vlan100(port1.0.11)
to not dynamic/not static
```

Support for IGMPv3/MLDv2

AlliedWare Plus supports IGMP version 3 and MLD version 2. What these versions add to the IGMP and MLD protocols is support for source-filtering. This means it now supports the ability for users to report interest in receiving a particular multicast stream only from a specific source or multiple source addresses. The desired source addresses can be expressed as either:

- Please send me the group A.B.C.D from any of the following source addresses (Include mode), or
- Please send me the group A.B.C.D from any source address except the following (Exclude mode)

This information can be used by multicast routing protocols to avoid delivering the multicast packets from specific sources to networks where there are no hosts interested in these multicast streams. This feature is also called Source Specific Multicast (SSM).

AlliedWare Plus IGMPv3 is based on the IETF RFC 3376 and the MLDv2 implementation is based on RFC3810. IGMPv3/MLDv2 supports just two types of protocol messages:

- Membership query messages
- Membership report messages

So, Leave/Done messages are no longer meaningful in these protocol versions; they are treated as such special cases of Reports.

An IGMPv3/MLDv2 Querier can query three types of information using a query message:

- A general query to determine the multicast reception state (membership information) associated with its neighboring hosts' interfaces for all multicast groups enabled at the interface.
- A group-specific query to determine the multicast reception state for a specific multicast group associated with its neighboring host's interfaces.
- A group-and-source-specific query to determine the multicast reception state for a specific multicast group and a specified set of sources, associated with its neighboring host's interfaces.

The first two types of query are equivalent to IGMPv2/MLDv1 queries. But the third type of query (the group-and-source-specific query) is a new type that is introduced by IGMPv3/MLDv2.

In IGMPv3/MLDv2, when a client has sent a general query, it is sent with an IP destination address of 224.0.0.1— this is the all-systems multicast address. Group-Specific and Group-and-Source-Specific Queries are sent with an IP destination address equal to the multicast address of interest.

As mentioned above, we need to point out that there is no such thing as an IGMPv3/MLDv2 Leave/Done message. A host leaves by requesting to receive a group from no sources.

A single IGMPv3/MLDv2 report can contain multiple Group Records. What this means is that the report can update a host's membership of multiple multicast groups. When the IGMPv3 reports are sent from the host, they are sent with a destination address of 224.0.0.22. and MLDv2 reports are all sent to the IPv6 address FF02::16. This results in all IGMPv3/MLDv2 multicast-aware routers listening to the reports.

A system that is operating in IGMPv1 or IGMPv2 compatibility modes sends version 1 or version 2 reports to the multicast group specified in the Group Address field of the Report, similarly a system running in MLDv1 compatibility mode sends MLDv1 reports to the multicast group in the Group Address field.

An implementation of IGMPv3 must also support the following three message types, for inter-operation with previous versions of IGMP:

- Version 1 Membership Report [RFC1112]
- Version 2 Membership Report [RFC2236]
- Version 2 Leave Group [RFC2236]

An implementation of MLDv2 must also support the following message types, to interoperate with MLDv1:

- Version 1 Membership Report [RFC2710]
- Version 1 Listener Done Group [RFC2710]

Summary of MLDv2 Packet Types

Table 6: Summary of MLDv2 packet types

PACKET TYPE	SOURCE ADDRESS	DEST ADDRESS	ICMP TYPE	MULTICAST ADDRESS IN PACKET
General Query	Sender's link-local address.	FF02::1	130	::
Multicast address-specific Query	Sender's link-local address.	The specific group address being queried about.	130	The specific group address being queried about.
Multicast address-and- source-specific Query	Sender's link-local address.	The specific group address being queried about.	130	The specific group address being queried about.
Report	Sender's link-local address.	FF02::16	143	Multiple multicast addresses.

Benefits

The benefits of using IGMPv3/MLDv2 over earlier versions include:

- Optimized bandwidth utilization: a receiver may request to receive traffic only from explicitly known sources
- Improved security (no denial of service attacks from unknown sources)
- Simplified service: enables service providers to receive content from multiple content providers without needing to be concerned if the content providers are using overlapping group addresses

Summary of differences between IGMPv2/MLDv1 and IGMPv3/MLDv2

The fundamental change between IGMPv2/MLDv1 and IGMPv3/MLDv2 is the fact that IGMPv3/MLDv2 implements source filtering.

This has flow-on effects like:

- IGMPv3/MLDv2 queriers keep separate forwarding entries for each source from which it is forwarding a group G, and can delete individual (S,G) forwarding entries, while retaining other entries for forwarding G from other sources.
- IGMPv3/MLDv2 introduced Multicast-Address-And-Source-Specific Queries.
- IGMPv3/MLDv2 Reports have been fundamentally altered. Rather than containing just a multicast group that is being requested, they contain Multiple Address Records that each contain a set of information relating to their relationship to a Group G.
- IGMPv3/MLDv2 has done away with the Listener Done message, as the redesigned Report packet has subsumed any purpose that a Listener Done message would fulfil.

- Multiple Groups can be reported in the same IGMPv3/MLDv2 Report packet.
- MLDv2 reports have a new ICMP type value
- IGMPv3/MLDv2 reports are sent to a constant destination multicast address.
- The S flag has been introduced to the Query message, to deal with difficulties that can arise when non-queriers do not receive reports that the querier does receive.
- The querier puts the values of its Robustness Variable and Query Interval into its Queries, so that non-queriers can adopt these values, thereby enabling a smoother transition if the current querier is replaced by one of the non-queriers.
- The range of possible values of Maximum Response Times has been increased by using an algorithm that can encode a floating-point value into the Maximum Response Time field of Queries.
- IGMPv3/MLDv2 hosts never suppress their Reports. All hosts respond to all Queries to which they have a meaningful response.

Backward compatibility to IGMPv2/MLDv1 has been built into IGMPv3/MLDv2.

Note: IGMPv3/MLDv2 Multicast address records that are equivalent to an IGMPv2/MLDv1 Report and an IGMPv2/MLDv1 Listener Done are:

- An IGMPv2/MLDv1 Report is equivalent to a CHANGE_TO_EXCLUDE_MODE Multicast Address Record with an empty source list. Such a Multicast Address Record says "***I have started listening to G on this interface, and I will accept G from ANY source***". This is exactly what an IGMPv2/MLDv1 Report says.
- An IGMPv2Leave/MLDv1 Listener Done is equivalent to a CHANGE_TO_INCLUDE_MODE Multicast Address Record with an empty source list. Such a Multicast Address Record says "***I have been using Exclude mode source filtering for G on this interface, but now, I am changing to Include Mode, and actually do not want to accept G from ANY sources at all***." This has the same effect as an IGMPv2 Leave/MLDv1 Listener Done.

SSM Mapping

Source Specific Multicasting expects that hosts subscribing to a channel will use IGMPv3/MLDv2 to specify the IP address of the source from which they wish to receive the channel.

However, the fact is that there is a large installed base of equipment that supports on IGMPv2/MLDv1, and not IGMPv3/MLDv2. It would be extremely annoying to not be able to use SSM in a network simply because some of the equipment connected to it does not support IGMPv3/MLDv2.

Consider the case of a service providing delivering TV as multicast over Ethernet. If this provider is receiving content from upstream content providers who only support PIM SSM and will not accept any (*,G) Joins, then the service provider must implement PIM SSM in their network. However, it is highly likely that at least some of the subscribers connected to the network will be using Set Top Boxes that are not capable of IGMPv3/MLDv2. This service provider is then stuck between a rock and a hard place, they either need to go around and replace ALL subscribers' Set Top Boxes with IGMPv3/MLDv2 capable devices, or they need their content providers to relax their (S,G) Join requirements. Neither of these options is going to be easy.

Fortunately, there is a third option, the multicast routers in the network could help them out, and provide a work-around that converts IGMPv2/MLDv1 reports into Source-Specific reports.

This third option is exactly what AlliedWare Plus provides.

To configure this feature, proceed as follows:

1. Create an access-list to define a range of multicast group addresses.

```
access-list 10 permit 232.1.67.0 0.0.0.255
ipv6 access-list standard SSM-groups permit FF35::/32
```

2. Enable SSM mapping of IGMPv1/v2 and MLDv1 reports.

```
ip igmp ssm-map enable
ipv6 mld ssm-map enable
```

3. Create a mapping which informs the router that all group addresses matching the access list are deemed to come from a certain source address.

```
ip igmp ssm-map static 10 89.156.213.78
ipv6 mld ssm-map static SSM-groups FF0E::1/128
```

The effect of this command is that if the router receives any IGMPv1 or IGMPv2 reports for Groups in the 232.1.67.0/24 range, then it will treat those reports as though they were a Source-Specific request, and the requested source was 89.156.213.78.

Also, if the router receives MLDv1 reports for groups in the FF35::/32 range, then it will treat those reports as though they were a Source-Specific request, and the requested source was FF0E::1/128.

Report Suppression

Snoopers can apply some judgment to the Report and Leave/Done messages that they forward to the Querier.

For example, if a snooper is currently receiving a given stream, and another downstream host sends in a report requesting that same stream, then forwarding that report up to the Querier is of no value, as the Querier will not need to make any changes based on receiving that report.

Similarly, if a snooper is forwarding a given stream to multiple downstream hosts, and one of those hosts sends in a leave for that stream, the snooper does not need to forward the leave up to the querier, as the snooper needs to continue receiving the stream to deliver to the remaining listening hosts.

So, snoopers can apply some judgment as to which reports and Leave/Done messages they forward to the Querier, to save having the Querier waste effort on processing messages that it does not need to. This process of applying judgment, and only forwarding the IGMP/MLD packets that need to be forwarded, is called **report suppression**.

Report suppression does not apply to IGMPv3/MLDv2. This functionality is enabled by default for IGMPv1, IGMPv2, MLDv1.

The command below enables report suppression for IGMPv2 and MLDv1 reports on vlan3, if it has been disabled.

```
awplus# configure terminal
awplus(config)# int vlan3
awplus(config-if)# ip igmp version 2
awplus(config-if)# ip igmp snooping report-suppression
awplus(config-if)# ipv6 mld version 1
awplus(config-if)# ipv6 mld snooping report-suppression
```

When Client A in vlan3 sends a report to join the group 225.1.2.3, the switch sends this report upstream to the Querier. If Client B in the same VLAN sends a report for the same multicast group 225.1.2.3, the switch does not need to forward this report on to the Querier—it had already done so when receiving the report from Client A.

The switch does not drop the report. It processes the second report to update (refresh) the liveness timer of the group in its group database, and then not forward it.

When the first client sends a leave message, the switch does not pass the leave message on up to the Querier because report suppression is on. As a result, the client will continue to receive traffic until the snooping entry on the port times out, or until the Querier sends a query for this group and the client does not confirm its membership for the group.

When report suppression is enabled, the switch will put its own IP address as the source address of the reports and leave messages that it sends to the Querier.

When the last client in a group leaves, this causes traffic to stop to all clients straight away (even to the clients that have already left but are still receiving traffic).

This is simply because a Leave/Done report from the last group member will not be suppressed (regardless of the suppression setting), but will be immediately forwarded to the Querier, so the leave sequence is started immediately.

If report suppression is disabled on switches in the network, ensure that immediate leave is not configured on any switches upstream of those on which report suppression has been disabled. This is because when report suppression is enabled, the switch will only send a leave upstream if all of its downstream hosts have left the group, but when report suppression is disabled, all leave messages received from downstream hosts are simply passed on upstream. If a port upstream of a suppression-disabled switch is configured for immediate leave, it will stop forwarding the stream as soon as any downstream host leaves the group, irrespective of whether other downstream hosts are still members of the group.

Source Address Check

By default AlliedWare Plus only accepts IGMP packets whose source IP address is in the same subnet as the IP address on the receiving interface. We can override this behavior using the command:

```
Switch-2(config)# int vlan100
Switch-2(config-if)# no ip igmp source-address-check
```

Tips for Making an IGMP/MLD Network More Efficient

- Enable Fast Leave—this will cause IGMP/MLD snooping to immediately remove group membership of the port, rather than waiting till the Leave report has gone to the Querier and the Querier has sent a query in response, asking for group membership.
- Switch off Report Suppression—this will cause all reports and Leaves from clients to be immediately forwarded to the Querier, causing the leave sequence to commence straight away. This is OK if only a small to medium number of clients are connected. Generally, it is not recommended if there are a very large number of clients, as the Querier will see a lot more reports than if the edge switches aggregate the reports from the clients.
- Move the Querier to the edge switch—this will have the desired effect because the clients can talk directly with the Querier, however it is not recommended if you have more than one edge switch.
- Reduce the Query Interval on the Querier—this will cause the Querier to send queries more often and therefore will shorten the time between the client leaving and the general query arriving that will cause the edge switch to delete the group entry on the client port.

Tips for Large Multicast Networks

From software version 5.4.8-1.2 onwards:

- SBx908 GEN2 switches support PIM-SM networks of up to 1024 interfaces
- SBx908 GEN2 system and SBx8100 systems containing SBx81CFC960 and SBx81XLEM cards support 8k (8192) multicast groups

From software version 5.4.9-0.1 onwards:

- SBx908 GEN2 switches support 32k (32,768) multicast groups
- x950 Series switches support PIM-SM networks of up to 1024 interfaces

This section discusses commands that help increase performance and simplify management of large multicast networks on these switches.

IGMP Querier on SBx908 GEN2 switches

If the SwitchBlade x908 GEN2 is acting as an IGMP querier, and you have 8192 or more multicast groups, you must turn off IGMP report suppression on all VLANs requiring IGMP group joins. To do this, use the command:

```
awplus(config-if)#no ip igmp snooping report-suppression
```

Silicon profile 3 on SBx8100 Series switches

To support 8192 multicast groups on the SwitchBlade x8100, you also need to change to silicon-profile 3 with a routing ratio weighting of multicast. Use the commands:

```
awplus(config)#platform silicon-profile profile3
awplus(config)#platform routingratio ipv4andipv6 weighting multicast
```

Create only (S,G) entries in hardware

When there are many multicast groups with many downstream interfaces, it can be helpful to prevent a large number of (*,G) entries from being created in hardware, using the command:

```
awplus(config)#ip multicast-routing [vrf <vrf-name>] ssm-only-hw
```

This command suppresses the creation of the (*,G) entries in hardware, but does not suppress the (S,G) entries. This improves the performance. (*,G) entries will still be created as needed in the CPU.

However, using this command may cause multicast data to be briefly flooded on the incoming interface if it comes from interfaces other than the “correct” incoming interface. The “correct” incoming interface is determined by Unicast Reverse Path Forwarding (uRPF).

Reduction of traffic to CPU

As part of software version 5.4.8-2.3, AlliedWare Plus now has the ability to block packets from coming up to the CPU when the multicast packet is seen on the wrong interface. By default this option is off and can be turned on using the following command.

```
awplus(config)#ip pim sparse-mode wrong-vif suppression
```

This will turn on the ability to block multicast packets from coming up to the CPU when they are on the wrong interface. This is done by creating a hardware entry for this multicast stream and dropping the packet instead of bringing it up to the CPU. NOTE: For each multicast stream seen on the wrong interface it will consume an entry normally used for multicast routing. For example if the network has 10K multicast streams on the wrong interface, this leaves only 22K hardware entries left for normal multicast traffic on a SBx908 GEN2.

Join/Prune message batching

From software version 5.4.8-2.6 onwards, IPv4 PIM-SM can be configured to attempt to batch Join and Prune messages for multiple groups in one PIM message packet. By default this is disabled and can be enabled with the following command.

```
awplus(config)#ip pim sparse-mode join-prune-batching
```

This will enable batching of PIM-SM Join and Prune messages, reducing the number of packets sent by the PIM-SM daemon when many groups are in use, it is recommended to enable this option if more than 4K streams are in use.

Filtering show command output

If you want to see information about a single interface or group, you can use the | symbol and filter the output. There are two particularly useful options: **begin** and **include**.

| begin This skips all the output before the first line that has the specified text. For example, to see the PIM interface details of vlan555 (and onwards), use the command:

```
awplus#show ip pim sparse-mode interface detail | begin vlan555
```

This gives the following output.

```
awplus#show ip pim sparse-mode interface detail | begin vlan555
...skipping
vlan555 (vif 555):
  Address 192.168.1.149, DR 192.168.1.149
  Hello period 30 seconds, Next Hello in 15 seconds
  Triggered Hello period 5 seconds
  Neighbors:
   192.168.1.22
```

| include This only displays lines of output that contain the specified text. For example, to see which VLAN has the address 10.2.104.10, use the command:

```
awplus#show ip pim sparse-mode interface detail | include 10.2.104.10
```

This gives the following output.

```
awplus#show ip pim sparse-mode interface | include 10.2.104.10
10.2.104.10      vlan1004  8        v2/S   0      1
10.2.104.10
```

Viewing the maximum number of multicast groups supported in hardware

You can see the maximum number of multicast groups by checking the route limit/route threshold setting in the **show ip mroute count** command:

```
awplus#show ip mroute count

IP Multicast Statistics
Total 8 routes using 1408 bytes memory
Route limit/Route threshold: 32768/32768
Total NOCACHE/WRONGVIF/WHOLEPKT recv from fwd: 398/0/0
Total NOCACHE/WRONGVIF/WHOLEPKT sent to clients: 398/0/0
Immediate/Timed stat updates sent to clients: 0/80
Reg ACK recv/Reg NACK recv/Reg pkt sent: 0/0/0
Next stats poll: 00:01:04

Forwarding Counts: Pkt count/Byte count, Other Counts: Wrong If pkts
Fwd msg counts: WRONGVIF/WHOLEPKT recv
Client msg counts: WRONGVIF/WHOLEPKT/Imm Stat/Timed Stat sent
Reg pkt counts: Reg ACK recv/Reg NACK recv/Reg pkt sent

(192.168.2.121, 224.6.1.1), Forwarding: 11/10, Other: 0
  Fwd msg: 0/0, Client msg: 0/0/0/10,  Reg: 0/0/0

(192.168.2.121, 224.6.1.2), Forwarding: 11/10, Other: 0
  Fwd msg: 0/0, Client msg: 0/0/0/10,  Reg: 0/0/0

(192.168.2.121, 224.6.1.3), Forwarding: 11/10, Other: 0
  Fwd msg: 0/0, Client msg: 0/0/0/10,  Reg: 0/0/0

(192.168.2.121, 224.6.1.4), Forwarding: 11/10, Other: 0
  Fwd msg: 0/0, Client msg: 0/0/0/10,  Reg: 0/0/0

(192.168.2.121, 224.6.1.5), Forwarding: 11/10, Other: 0
  Fwd msg: 0/0, Client msg: 0/0/0/10,  Reg: 0/0/0

(192.168.2.121, 224.6.1.6), Forwarding: 11/10, Other: 0
  Fwd msg: 0/0, Client msg: 0/0/0/10,  Reg: 0/0/0

(192.168.2.121, 224.6.1.7), Forwarding: 11/10, Other: 0
  Fwd msg: 0/0, Client msg: 0/0/0/10,  Reg: 0/0/0

(192.168.2.121, 224.6.1.8), Forwarding: 11/10, Other: 0
  Fwd msg: 0/0, Client msg: 0/0/0/10,  Reg: 0/0/0
```

When 32,000 entries are learnt the show command displays can get very long and take a long time to complete. Often administrator are only interested in the total number of entries learnt, so using the ‘brief’ version of show commands can be useful:

To display the L2 IGMP groups learnt and the number of interfaces configured use the command:

```
awplus# show ip igmp groups brief
```

```
awplus#show ip igmp groups brief
IGMP Groups Brief

  IGMP Configured Interfaces             98
  IGMP Configured Interfaces (Up)        98
  IGMP Configured Interfaces (Down)      0

  IGMP Stopped Groups                    32000
  IGMP Dynamic Groups                    32000
  IGMP Static Groups                     0
```

To display L3 PIM sparse-mode multicast entries learnt, use the command:

```
awplus# show ip pim sparse-mode mroute brief
```

```
awplus#show ip pim sparse-mode mroute brief
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 0
(S,G) Entries: 32000
(S,G,rpt) Entries: 32000
FCR Entries: 0
MRIB Msg Cache Hit: 0
```

To see the number of multicast entries learnt in hardware and which downstream interfaces are used, use the command:

```
awplus# show platform table ipmulti-brief.
```

Stack members should contain the same number of entries across all members. The **Hit** bit parameter shows the route is still active. Normally, you would only expect it to be active on one stack member, but there are cases where it should be present on both stack members. If there are no port counts there are no downstream members and no traffic will be forwarded. The **show platform table ipmulti** command can be used to debug this further.

```
awplus# sh platform table ipmulti-brief

Stack member 1:

[Instance 4]
Source Ip         Multicast         VLAN   L2 Port   L3 Port   Hit
Addr              IP Addr           ID     Count     Count     Bit
-------------------------------------------------------------------
10.36.20.1        239.255.1.5       10     2         1         yes
10.36.20.1        239.255.1.4       10     2         1         yes
10.36.20.1        239.255.1.3       10     2         1         yes
10.36.20.1        239.255.1.2       10     2         1         yes
10.36.20.1        239.255.1.1       10     2         1         yes
2001:db8:ffff::1  ff08::5           10     3         1          no
2001:db8:ffff::1  ff08::4           10     3         1          no
2001:db8:ffff::1  ff08::2           10     3         1          no
2001:db8:ffff::1  ff08::3           10     3         1          no


Stack member 2:

[Instance 8]
Source Ip         Multicast         VLAN   L2 Port   L3 Port   Hit
Addr              IP Addr           ID     Count     Count     Bit
-------------------------------------------------------------------
10.36.20.1        239.255.1.5       10     2         1          no
10.36.20.1        239.255.1.4       10     2         1          no
10.36.20.1        239.255.1.3       10     2         1          no
10.36.20.1        239.255.1.2       10     2         1          no
10.36.20.1        239.255.1.1       10     2         1          no
2001:db8:ffff::1  ff08::5           10     3         1         yes
2001:db8:ffff::1  ff08::4           10     3         1         yes
2001:db8:ffff::1  ff08::2           10     3         1         yes
2001:db8:ffff::1  ff08::3           10     3         1         yes
```