7

Multicast Listener Discovery (MLD) Snooping

Contents

Overview

Multicast addressing allows one-to-many or many-to-many communication among hosts on a network. Typical applications of multicast communication include audio and video streaming, desktop conferencing, collaborative computing, and similar applications.

Multicast Listener Discovery (MLD) is an IPv6 protocol used on a local link for multicast group management. MLD is enabled per VLAN, and is analogous to the IPv4 IGMP protocol.

MLD snooping is a subset of the MLD protocol that operates at the port level and conserves network bandwidth by reducing the flooding of multicast IPv6 packets.

This chapter describes concepts of MLD snooping and the CLI commands available for configuring it and for viewing its status.

Introduction to MLD Snooping

There are several roles that network devices may play in an IPv6 multicast environment:

- **MLD host**—a network node that uses MLD to "join" (subscribe to) one or more multicast groups
- **multicast router**—a router that routes multicast traffic between subnets
- **querier**—a switch or multicast router that identifies MLD hosts by sending out MLD queries, to which the MLD hosts respond

Curiously enough, a network node that acts as a *source* of IPv6 multicast traffic is only an indirect participant in MLD snooping—it just provides multicast traffic, and MLD doesn't interact with it. (Note, however, that in an application like desktop conferencing a network node may act as both a source and an MLD host; but MLD interacts with that node only in its role as an MLD host.)

A source node creates multicast traffic by sending packets to a multicast address. In IPv6, addresses with the first eight bits set (that is, "FF" as the first two characters of the address) are multicast addresses, and any node that listens to such an address will receive the traffic sent to that address. Application software running on the source and destination systems cooperates to determine what multicast address to use. (Note that this is a function of the application software, not of MLD.)

For example, if several employees engage in a desktop conference across the network, they all need application software on their computers. At the start of the conference, the software on all the computers determines a multicast address of, say, FF3E:30:2001:DB8::101 for the conference. Then any traffic sent to that address can be received by all computers listening on that address.

General operation. Multicast communication can take place without MLD, and by default MLD is disabled. In that case, if a switch receives a packet with a multicast destination address, it floods the packet to all ports in the same VLAN (except the port that it came in on). Any network nodes that are listening to that muticast address will see the packet; all other hosts ignore the packet.

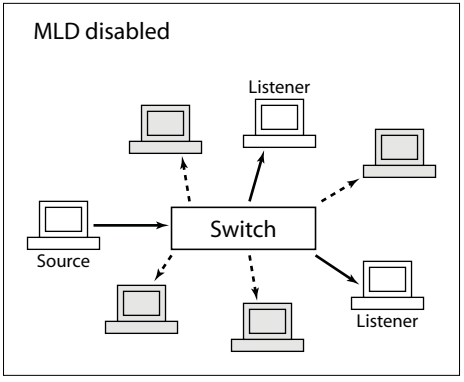

Figure 7-1. Without MLD, multicast traffic is flooded to all ports.

When MLD snooping is enabled on a VLAN, the switch acts to minimize unnecessary multicast traffic. If the switch receives multicast traffic destined for a given multicast address, it forwards that traffic only to ports on the VLAN that have MLD hosts for that address. It drops that traffic for ports on the VLAN that have no MLD hosts (except for a few special cases explained below).

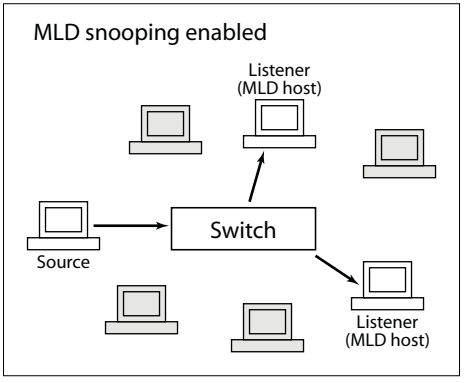

Figure 7-2. With MLD snooping, traffic is sent to MLD hosts.

Note that MLD snooping operates on a single VLAN (though there can be multiple VLANs, each running MLD snooping). Cross-VLAN traffic is handled by a multicast router.

Forwarding in MLD snooping. When MLD snooping is active, a multicast packet is handled by the switch as follows:

- forwarded to ports that have nodes that have joined the packet's multicast address (that is, MLD hosts on that address)
- forwarded toward the querier—If the switch is not the querier, the packet is forwarded out the port that leads to the querier.
- forwarded toward any multicast routers—If there are multicast routers on the VLAN, the packet is forwarded out any port that leads to a router.
- forwarded out administratively forwarded ports—The packet will be forwarded through all ports set administratively to forward mode. (See the description of forwarding modes, below.)
- dropped for all other ports

Each individual port's forwarding behavior can be explicitly set using a CLI command to one of these modes:

- auto (the default mode)—The switch forwards packets through this port based on the MLD rules and the packet's multicast address. In most cases, this means that the switch forwards the packet only if the port connects to a node that is joined to the packet's multicast address (that is, to an MLD host). There is seldom any reason to use a mode other than "auto" in normal operation (though some diagnostics may make use of "forward" or "block" mode).
- forward—The switch forwards all IPv6 multicast packets through the port. This includes IPv6 multicast data and MLD protocol packets.
- block—The switch drops all MLD packets received by the port and blocks all outgoing IPv6 multicast packets through the port, except those packets destined for well known IPv6 multicast addresses. This has the effect of preventing IPv6 multicast traffic from moving through the port.

Note that the switch floods all packets with "well known" IPv6 multicast destination addresses through all ports. Well known addresses are permanent addresses defined by the Internet Assigned Numbers Authority (www.iana.org). IPv6 standards define any address beginning with FF0x/12 (binary 1111 1111 0000) as a well known address.

Listeners and joins. The "snooping" part of MLD snooping arises because a switch must keep track of which ports have network nodes that are MLD hosts for any given multicast address. It does this by keeping track of "joins" on a per-port basis.

A network node establishes itself as an MLD host by issuing a multicast "join" request (also called a multicast "report") for a specific multicast address when it starts an application that listens to multicast traffic. The switch to which the node is connected sees the join request and forwards traffic for that multicast address to the node's port.

Queries. The querier is a multicast router or a switch that periodically asks MLD hosts on the network to verify their multicast join requests. There is one querier for each VLAN, and all switches on the VLAN listen to the responses of MLD hosts to multicast queries, and forward or block multicast traffic accordingly.

All of the ProCurve switches described by this guide have the querier function enabled by default. If there is another device on the VLAN that is already acting as querier, the switch defers to that querier. If there is no device acting as querier, the switch enters an election state and negotiates with other devices on the network (if any) to determine which one will act as the querier.

The querier periodically sends general queries to MLD hosts on each multicast address that is active on the VLAN. The time period that the querier waits between sending general queries is known as the query interval; the MLD standard sets the default query interval to 125 seconds.

Network nodes that wish to remain active as MLD hosts respond to the queries with join requests; in this way they continue to assert their presence as MLD hosts. The switch through which any given MLD host connects to the VLAN sees the join requests and continues forwarding traffic for that multicast address to the MLD host's port.

Leaves. A node acting as an MLD host can be disconnected from a multicast address in two ways:

- It can stop sending join requests to the querier. This might happen if the multicast application quits or the node is removed from the network. If the switch goes for slightly more than two query intervals without seeing a join request from the MLD host, it stops sending multicast traffic for that multicast address to the MLD host's port.
- It can issue a "leave" request. This is done by the application software running on the MLD host. If the MLD host is the only node connected to its switch port, the switch sees the leave request and stops sending multicast packets for that multicast address to that port. (If there is more than one node connected to the port the situation is somewhat more complicated, as explained below under "Fast leaves and forced fast leaves".)

Fast leaves and forced fast leaves. The fast leave and forced fast leave functions can help to prune unnecessary multicast traffic when an MLD host issues a leave request from a multicast address. Fast leave is enabled by default and forced fast leave is disabled by default. Both functions are applied to individual ports.

Which function to use depends on whether a port has more than one node attached to it, as follows:

- If a port has only one node attached to it, then when the switch sees a leave request from that node (an MLD host) it knows that it does not need to send any more multicast traffic for that multicast address to the host's port. If fast leave is enabled (the default setting), the switch stops sending the multicast traffic immediately. If fast leave is disabled, the switch continues to look for join requests from the host in response to group-specific queries sent to the port. The interval during which the switch looks for join requests is brief and depends on the forced fast leave setting: if forced fast leave is enabled for the port, it is equal to the "forced fast leave interval" (typically a couple of seconds or less); if forced fast leave is disabled for the port, the period is about 10 seconds (governed by the MLD standard). When this process has completed the multicast traffic for the group will be stopped (unless the switch sees a new join request).
- If there are multiple nodes attached to a single port, then a leave request from one of those nodes (an MLD host) does not provide enough information for the switch to stop sending multicast traffic to the port. In this situation the fast leave function does not operate. The switch continues to look for join requests from any MLD hosts connected to the port, in response to group-specific queries sent to the port. As in the case described above for a single-node port that is not enabled for fast leave, the interval during which the switch looks for join requests is brief and depends on the forced fast leave setting. If forced fast leave is enabled for the port, it is equal to the "forced fast leave interval" (typically a couple of seconds or less); if forced fast leave is disabled for the port, the period is about 10 seconds (governed by the MLD standard). When this process has completed the multicast traffic for the group will be stopped unless the switch sees a new join request. This reduces the number of multicast packets forwarded unnecessarily.

Configuring MLD

Several CLI commands are available for configuring MLD parameters on a switch.

Enabling or Disabling MLD Snooping on a VLAN

Syntax: [no] ipv6 mld

Note: This command must be issued in a VLAN context.

This command enables MLD snooping on a VLAN. Enabling MLD snooping applies the last-saved or the default MLD configuration, whichever was most recently set.

The **[no]** *form of the command disables MLD snooping on a VLAN.*

MLD snooping is disabled by default.

For example, to enable MLD snooping on VLAN 8:

```
ProCurve# config
ProCurve(config)# vlan 8
ProCurve(vlan-8)# ipv6 mld
```

To disable MLD snooping on VLAN 8:

```
ProCurve(vlan-8)# no ipv6 mld
```

Configuring Per-Port MLD Traffic Filters

Syntax: ipv6 mld [auto <*port-list*> | blocked <*port-list*> | forward <*port-list*>]

Note: *This command must be issued in a VLAN context.*

This command sets per-port traffic filters, which specify how each port should handle MLD traffic. Allowed settings are:

auto—*follows MLD snooping rules: packets are forwarded for joined groups*

blocked—*all multicast packets are dropped, except that packets for well known addresses are forwarded*

forward—*all multicast packets are forwarded*

The default value of the filter is **auto**.

<*port-list*>—*specifies the affected port or range of ports*

For example:

```
ProCurve(vlan-8)# ipv6 mld forward a16-a18
ProCurve(vlan-8)# ipv6 mld blocked a19-a21
ProCurve(vlan-8)# show ipv6 mld vlan 8 config

 MLD Service Vlan Config

  VLAN ID : 8
  VLAN NAME : VLAN8
  MLD Enabled [No] : Yes
  Querier Allowed [Yes] : Yes

  Port Type      | Port Mode Forced Fast Leave Fast Leave
  ---- --------- + --------- ----------------- ----------
  A13  100/1000T | auto      No                Yes
  A14  100/1000T | auto      No                Yes
  A15  100/1000T | auto      No                Yes
  A16  100/1000T | forward   No                Yes
  A17  100/1000T | forward   No                Yes
  A18  100/1000T | forward   No                Yes
  A19  100/1000T | blocked   No                Yes
  A20  100/1000T | blocked   No                Yes
  A21  100/1000T | blocked   No                Yes
  A22  100/1000T | auto      No                Yes
  A23  100/1000T | auto      No                Yes
  A24  100/1000T | auto      No                Yes
```

Figure 7-3. Example of an MLD Configuration with Traffic Filters

Configuring the Querier

Syntax: [no] ipv6 mld querier

Note: *This command must be issued in a VLAN context.*

This command enables the switch to act as querier on a VLAN.

The **[no]** *form of the command disables the switch from acting as querier on a VLAN.*

The querier function is enabled by default. If another switch or a multicast router is acting as the MLD querier on the VLAN, this switch will defer to that device. If an acting querier stops performing the querier function, all querier-enabled switches and multicast routers on the VLAN will enter an election to determine the next device to act as querier.

For example, to disable the switch from acting as querier on VLAN 8:

```
ProCurve(vlan-8)# no ipv6 mld querier
```

To enable the switch to act as querier on VLAN 8:

```
ProCurve(vlan-8)# ipv6 mld querier
```

Configuring Fast Leave

Syntax: [no] ipv6 mld fastleave <*port-list*>

Note: This command must be issued in a VLAN context.

This command enables the fast leave function on the specified ports in a VLAN.

The **[no]** *form of the command disables the fast leave function on the specified ports in a VLAN.*

The fast leave function is enabled by default.

For example, to disable fast leave on ports in VLAN 8:

```
ProCurve(vlan-8)# no ipv6 mld fastleave a14-a15
```

To enable fast leave on ports in VLAN 8:

```
ProCurve(vlan-8)# ipv6 mld fastleave a14-a15
```

Configuring Forced Fast Leave

Syntax: [no] ipv6 mld forcedfastleave <*port-list*>

Note: *This command must be issued in a VLAN context.*

This command enables the forced fast leave function on the specified ports in a VLAN.

The **[no]** *form of the command disables the forced fast leave function on the specified ports in a VLAN.*

The forced fast leave function is disabled by default.

For example, to enable forced fast leave on ports in VLAN 8:

```
ProCurve(vlan-8)# ipv6 mld forcedfastleave a19-a20
```

To disable forced fast leave on ports in VLAN 8:

```
ProCurve(vlan-8)# no ipv6 mld forcedfastleave a19-a20
```

Displaying MLD Status and Configuration

Current MLD Status

Syntax: show ipv6 mld

Displays MLD status information for all VLANs on the switch that have MLD configured.

show ipv6 mld vlan <vid>

Displays MLD status for the specified VLAN
vid—*VLAN ID*

For example, a switch with MLD snooping configured on VLANs 8 and 9 might show the following information:

```
ProCurve# show ipv6 mld

  MLD Service Protocol Info

  Total vlans with MLD enabled                : 2
  Current count of multicast groups joined    : 37

  VLAN ID : 8
  VLAN NAME : VLAN8
  Querier Address : fe80::218:71ff:fec4:2f00 [this switch]
  Querier Up Time : 1h:37m:20s
  Querier Expiry Time : 0h:1m:44s

  Ports with multicast routers :


  Active Group Addresses                   Type ExpiryTime Ports
  ---------------------------------------- ---- ---------- --------------------
  ff02::c                                  FILT 0h:4m:9s   A15-A21
  ff02::1:2                                FILT 0h:4m:3s   A21
  ff02::1:3                                FILT 0h:4m:9s   A15-A21
  ff02::1:ff00:42                          FILT 0h:4m:0s   A19
  ff02::1:ff02:2                           FILT 0h:4m:2s   A15
  ff02::1:ff02:3                           FILT 0h:4m:5s   A16
  ff02::1:ff03:2                           FILT 0h:4m:2s   A17
  ff02::1:ff03:3                           FILT 0h:4m:5s   A18
```

Figure 7-4. Example of Displaying the MLD Configuration for All Static VLANs on the Switch

```
 ff02::1:ff04:3                          FILT 0h:4m:5s   A20
 ff02::1:ff05:1                          FILT 0h:4m:3s   A21
 ff02::1:ff0b:2dfe                       FILT 0h:3m:59s  A17
 ff02::1:ff0b:d7d9                       FILT 0h:4m:4s   A15
 ff02::1:ff0b:da09                       FILT 0h:4m:5s   A18
 ff02::1:ff0b:dc38                       FILT 0h:4m:3s   A19
 ff02::1:ff0b:dc8d                       FILT 0h:4m:4s   A20
 ff02::1:ff0b:dd56                       FILT 0h:4m:0s   A16
 ff02::1:ff12:e0cd                       FILT 0h:4m:5s   A21
 ff02::1:ff4e:98a5                       FILT 0h:4m:0s   A17
 ff02::1:ff57:21a1                       FILT 0h:3m:58s  A20
 ff02::1:ff6b:dd51                       FILT 0h:4m:0s   A15
 ff02::1:ff7b:ac55                       FILT 0h:4m:5s   A16
 ff02::1:ff8f:61ea                       FILT 0h:4m:1s   A19
 ff02::1:ffc8:397b                       FILT 0h:4m:0s   A18
ff3e:30:2001:db8:8:0:7:101               FILT 0h:4m:4s   A15,A18,A21
 ff3e:30:2001:db8:8:0:7:102              FILT 0h:4m:13s  A16,A19

 VLAN ID : 9
 VLAN NAME : VLAN9
 Querier Address : fe80::218:71ff:fec4:2f00 [this switch]
 Querier Up Time : 1h:37m:22s
 Querier Expiry Time : 0h:1m:43s

Ports with multicast routers :

 Active Group Addresses                  Type ExpiryTime Ports
 --------------------------------------- ---- ---------- --------------------
 ff02::c                                 FILT 0h:4m:12s  B3,B5,B7
 ff02::1:3                               FILT 0h:4m:12s  B3,B5,B7
 ff02::1:ff02:4                          FILT 0h:4m:4s   B3
 ff02::1:ff03:4                          FILT 0h:3m:59s  B5
 ff02::1:ff04:4                          FILT 0h:4m:12s  B7
 ff02::1:ff0b:dc64                       FILT 0h:4m:0s   B7
 ff02::1:ff0b:dcf3                       FILT 0h:4m:2s   B3
 ff02::1:ff0b:dd5c                       FILT 0h:4m:4s   B5
 ff02::1:ff34:a69e                       FILT 0h:4m:1s   B5
 ff02::1:ff8e:11d5                       FILT 0h:3m:57s  B7
 ff02::1:ffea:2c4f                       FILT 0h:3m:58s  B3
```

Figure 7-5. Continuation of Figure 7-4

The following information is shown for each VLAN that has MLD snooping enabled:

- VLAN ID number and name
- Querier address: IPv6 address of the device acting as querier for the VLAN
- Querier up time: the length of time in seconds that the querier has been acting as querier
- Querier expiry time: If this switch is the querier, this is the amount of time until the switch sends the next general query. If this switch is not the querier, this is the amount of time in seconds until the current querier is considered inactive (after which a new querier election is held).
- Ports with multicast routers: ports on the VLAN that lead toward multicast routers (if any)
- Multicast group address information for each active group on the VLAN, including:
  - the multicast group address
  - the type of tracking for multicast joins: standard or filtered. If MLD snooping is enabled, port-level tracking results in filtered groups. If MLD snooping is not enabled, joins result in standard groups being tracked by this device. In addition, if hardware resources for multicast filtering are exhausted, new joins may result in standard groups even though MLD snooping is enabled.
  - expiry time: the time until the group expires if no joins are seen
  - the ports that have joined the multicast group

The group addresses you see listed typically result from several network functions. In our example, several of the addresses at the top of the list for each VLAN are IANA well known addresses (see www.iana.org/assignments/ipv6-multicast-addresses); the addresses in the form of ff02::1:ffxx:xxxx are solicited-node multicast addresses (used in IPv6 Neighbor Discovery); and the addresses beginning with ff3e are group addresses used by listeners to streaming video feeds.

Current MLD Configuration

Syntax: show ipv6 mld config

Displays current global MLD configuration for all MLD-enabled VLANS on the switch.

show ipv6 vlan *<vid>* config

Displays current MLD configuration for the specified VLAN, including per-port configuration information.
vid—*VLAN ID*

For example, the general form of the command might look like this:

```
ProCurve# show ipv6 mld config

 MLD Service Config

  Control unknown multicast  [Yes] : Yes
  Forced fast leave timeout  [4] : 4

  VLAN ID VLAN NAME       MLD Enabled Querier Allowed
  ------- --------------- ----------- ---------------
  8       VLAN8           Yes         Yes
  9       VLAN9           Yes         Yes
```

Figure 7-6. Example of a Global MLD Configuration

The following information, for all MLD-enabled VLANs, is shown:

- Control unknown multicast: If this is set to YES, any IPv6 multicast packets that are not joined by an MLD host will be sent only to ports that have detected a multicast router or ports that are administratively forwarded. If this is set to NO (or if MLD snooping is disabled), unjoined IPv6 multicast packets will be flooded out all ports in the VLAN.
- Forced fast leave timeout: the interval between an address specific query and a forced fast leave (assuming no response), in tenths of seconds
- For each VLAN that has MLD enabled:
  - VLAN ID and name
  - whether MLD is enabled on the VLAN (default NO, but the VLAN will not show up on this list unless MLD is enabled)
  - whether the switch can act as querier for the VLAN (default YES)

The specific form of the command might look like this:

```
ProCurve# show ipv6 mld vlan 8 config

 MLD Service Vlan Config

  VLAN ID : 8
  VLAN NAME : VLAN8
  MLD Enabled [No] : Yes
  Querier Allowed [Yes] : Yes

  Port Type      | Port Mode Forced Fast Leave Fast Leave
  ---- --------- + --------- ----------------- ----------
  A13  100/1000T | auto      No                Yes
  A14  100/1000T | auto      No                Yes
  A15  100/1000T | auto      No                Yes
  A16  100/1000T | auto      No                Yes
  A17  100/1000T | auto      No                Yes
  A18  100/1000T | auto      No                Yes
  A19  100/1000T | auto      No                Yes
  A20  100/1000T | auto      No                Yes
  A21  100/1000T | auto      No                Yes
  A22  100/1000T | auto      No                Yes
  A23  100/1000T | auto      No                Yes
  A24  100/1000T | auto      No                Yes
```

Figure 7-7. Example of an MLD Configuration for a Specific VLAN

The following information is shown, if the specified VLAN is MLD-enabled:

- VLAN ID and name
- whether MLD is enabled on the VLAN (default NO, but the information for this VLAN will be listed only if MLD is enabled)
- whether the switch is allowed to act as querier on the VLAN

Ports Currently Joined

Syntax: show ipv6 vlan *<vid>* group

Lists the ports currently joined for all IPv6 multicast group addresses in the specified VLAN
vid—*VLAN ID*

show ipv6 vlan *<vid>* group *<ipv6-addr>*

Lists the ports currently joined for the specified IPv6 multicast group address in the specified VLAN
vid—*VLAN ID*
ipv6-addr—*address of the IPv6 multicast group for which you want information*

For example, the general form of the command is shown below. The specific form the the command is similar, except that it lists the port information for only the specified group.

```
ProCurve# show ipv6 mld vlan 9 group

 MLD Service Protocol Group Info

 VLAN ID : 9
 VLAN Name : VLAN9

  Filtered Group Address : ff02::c
 Last Reporter : fe80::7061:4b38:dbea:2c4f
  ExpiryTime : 0h:2m:19s

  Port  Port Type  | Port Mode  ExpiryTime
  ----  --------- + ---------  --------------------
  B3    100/1000T | auto       0h:2m:19s
  B5    100/1000T | auto       0h:2m:18s

    .
    .
    .

 Filtered Group Address : ff3e:30:2001:db8:9:0:7:111
  Last Reporter : fe80::7061:4b38:dbea:2c4f
  ExpiryTime : 0h:4m:14s

  Port  Port Type  | Port Mode  ExpiryTime
  ----  --------- + ---------  --------------------
  B3    100/1000T | auto       0h:4m:14s
  B5    100/1000T | auto       0h:4m:09s
```

Figure 7-8. Example of Ports Joined to Multicast Groups in a Specific VLAN

The following information is shown:

- VLAN ID and name
- port information for each IPv6 multicast group address in the VLAN (general group command) or for the specified IPv6 multicast group address (specific group command):
  - group multicast address
  - last reporter: last MLD host to send a join to the group address
  - group expiry time: the time until the group expires if no further joins are seen
  - port name for each port
  - port type for each port: Ethernet connection type
  - port mode for each port: auto (follows MLD snooping rules; that is, packets are forwarded for joined groups), forward (all multicast packets are forwarded to this group), or blocked (all multicast packets are dropped, except that packets for well-known addresses are forwarded)
  - expiry time for each port: amount of time until this port is aged out of the multicast address group, unless a join is received

Statistics

Syntax: show ipv6 mld statistics

Shows MLD statistics for all MLD-enabled VLANs

Syntax: show ipv6 mld vlan <vid> statistics

Shows MLD statistics for the specified VLAN
vid—*VLAN ID*

The general form the of the command shows the total number of MLD-enabled VLANs and a count of multicast groups currently joined. Both forms of the command show VLAN IDs and names, as well as the number of filtered and standard multicast groups and the total number of multicast groups.

For example, the general form of the command:

```
ProCurve# show ipv6 mld statistics

 MLD Service Statistics

  Total vlans with MLD enabled                : 2
  Current count of multicast groups joined    : 36


 MLD Joined Groups Statistics

  VLAN ID VLAN NAME    filtered     standard     total
  ------- ------------ ------------ ------------ ------------
  8       VLAN8        26           0            26
  9       VLAN9        10           0            10
```

Figure 7-9. Example of MLD Statistics for All VLANs Configured

And the specific form of the command:

```
ProCurve# show ipv6 mld vlan 8 statistics

 MLD Statistics

  VLAN ID : 8
  VLAN NAME : VLAN8

  Number of Filtered Groups     : 26
  Number of Standard Groups     : 0
  Total Multicast Groups Joined : 26
```

Figure 7-10. Example of MLD Statistics for a Single VLAN

Counters

Syntax: show ipv6 mld vlan *<vid>* counters

Displays MLD counters for the specified VLAN
vid—*VLAN ID*

```
ProCurve# show ipv6 mld vlan 8 counters

 MLD Service Vlan Counters

  VLAN ID : 8
  VLAN NAME : VLAN8

     General Query Rx                   : 2
     General Query Tx                   : 0
     Group Specific Query Rx            : 0
     Group Specific Query Tx            : 0
     V1 Member Report Rx                : 1589
     V2 Member Report Rx                : 15
     Leave Rx                           : 30
     Unknown MLD Type Rx                : 0
     Unknown Pkt Rx                     : 0
     Forward to Routers Tx Counter      : 83
     Forward to Vlan Tx Counter         : 48
     Port Fast Leave Counter            : 4
     Port Forced Fast Leave Counter     : 0
     Port Membership Timeount Counter   : 28
```

Figure 7-11. Example of MLD Counters for a Single VLAN

The following information is shown:

- VLAN number and name
- For each VLAN:
  - number of general queries received
  - number of general queries sent
  - number of group-specific queries received
  - number of group-specific queries sent
  - number of MLD version 1 member reports (joins) received
  - number of MLD version 2 member reports (joins) received
  - number of leaves received
  - number of MLD packets of unknown type received
  - number of packets of unknown type received
  - number of packets forwarded to routers on this VLAN
  - number of times a packet has been forwarded to all ports on this VLAN
  - number of fast leaves that have occurred
  - number of forced fast leaves thathave occurred
  - number of times a join has timed out on this VLAN