Chapter 11

Configuring Multicast Listener Discovery

Hosts use Multicast Listener Discovery (MLD) protocol in IPv6 to report their multicast group memberships to neighboring routers. Similarly, multicast routers, such as the E-series router, use MLD to discover which of their hosts belong to multicast groups.

This chapter describes how to configure MLD on an E-series router; it contains the following sections:

Overview

The IPv6 address scheme uses hexadecimal FF at the start of an address for IPv6 multicast. MLD is a protocol that uses these addresses. The following addresses have specific functions:

- You can assign only multicast addresses of global-scope (that is, containing an FF*x*E prefix, where *x* is the flags field) to a multicast group.
- FF02::1 is the link-scope all-nodes address—A packet sent to this address reaches all nodes on a subnetwork.
- FF02::2 is the link-scope all-routers address—A packet sent to this address reaches all routers on a subnetwork.
- FF02::16 is the link-scope all-MLDv2 routers address—A packet sent to this address reaches all MLDv2 routers on a subnetwork.

This implementation of MLD complies with MLD versions 1 and 2. MLDv2 allows for source-specific join and leave messages and is backward compatible with MLDv1. Configuring MLDv1 with the SSM mapping feature provides support for source-specific joins.

MLDv1 mode interfaces exchange the following types of messages between routers and hosts:

- Multicast listener queries
- Multicast listener reports
- Multicast listener done messages

MLDv2 mode interfaces exchange the following types of messages with MLDv2 hosts:

- Multicast listener queries
- MLDv2 multicast listener reports

Multicast Listener Queries

A multicast router can be a querier or a nonquerier. There is only one querier on a network at any time. Multicast routers monitor queries from other multicast routers to determine the status of the querier. If the querier hears a query from a router with a lower IPv6 address, it relinquishes its role to that router.

MLDv1 and MLDv2 mode interfaces send two types of multicast listener queries to hosts on the network:

- General queries to the all-nodes address (FF02::1)
- Specific queries to the appropriate multicast group address

MLDv2 mode interfaces send the following type of queries to MLDv2 hosts:

- General queries
- Group-specific queries
- Source-specific queries

The purpose of a membership group query is to discover the multicast groups to which a host belongs.

MLDv1 and MLDv2 multicast listener queries have a Max Response Time field. This response time is the maximum that a host can take to reply to a query.

Multicast Listener Reports

When a host receives a multicast listener query, it identifies the groups associated with the query and determines to which groups it belongs. The host then sets a timer, with a value less than the Max Response Time field in the query, for each group to which it belongs.

When the timer expires, the host sends a multicast listener report to the group address. When a multicast router receives a report, it adds the group to the membership list for the network and sets a timer to the *multicast address listening interval*. If this timer expires before the router receives another multicast listener report, the router determines that the group has no members left on the network.

If the router does not receive any reports for a specific multicast group within the *maximum response time*, it determines that the group has no members on the network. The router does not forward subsequent multicasts for that group to the network.

MLDv2 supports an extended report format that allows you to report multiple groups and source lists in a single report. These reports are addressed to the all-MLDv2 router's multicast address (FF02::16).

Multicast Listener Done Messages

When an MLDv1 host leaves a group, it sends a multicast listener done message to multicast routers on the network. A host generally addresses multicast listener done messages to the all-routers address, FF02::2.

When an MLDv2 host leaves a group, it sends a multicast listener report. This report includes an empty source list for that group.

Platform Considerations

For information about modules that support MLD on the ERX-7xx models, ERX-14xx models, and the ERX-310 router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support MLD.

For information about modules that support MLD on the E120 router and the E320 router:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support MLD.

References

For more information about MLD, see the following resources:

- RFC 3710—Multicast Listener Discovery (MLD) for IPv6 (October 1999) on page 578
- IGMP/MLD-based Multicast Forwarding ('IGMP/MLD Proxying')—draft-ietf-magma-igmp-proxy-06.txt (October 2004 expiration) on page 587
- Multicast Group Membership Discovery MIB—draft-ietf-magma-mgmd-mib-06.txt (October 2004 expiration) on page 587

Before You Begin

You can configure MLD only on IPv6 interfaces. For information about configuring IPv6 interfaces, see *JUNOSe IP, IPv6, and IGP Configuration Guide, Chapter 2, Configuring IPv6*.

Configuring Static and Dynamic MLD Interfaces

The router supports *static* and *dynamic* MLD interfaces. Unlike static interfaces, dynamic interfaces are not restored when you reboot the router. For some protocols, dynamic layers can build on static layers in an interface; however, in a dynamic MLD interface, all the layers are dynamic. See Figure 17 for examples of static and dynamic MLD interfaces.

Figure 17: Static and Dynamic MLD Interfaces

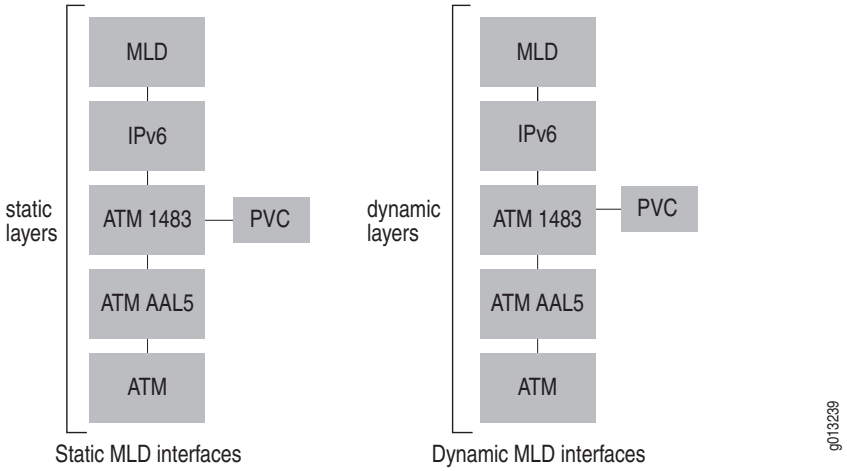

Static MLD interfaces are configured with software such as the CLI or an SNMP application; dynamic MLD interfaces are configured with a profile. A profile comprises a set of attributes for an interface; a profile for dynamic MLD interfaces contains attributes for configuring all the layers in the interface.

You define a profile by using the same CLI commands that you use to configure a static MLD interface; however, the mode in which you use the commands differs. Use the commands in Interface Configuration mode to configure a static MLD interface and in Profile Configuration mode to define a profile.

When you have defined a profile, you can apply it to an interface or a group of interfaces. Profiles provide an efficient method of creating and managing large numbers of dynamic interfaces. For detailed information about creating and assigning profiles, see *JUNOSe Link Layer Configuration Guide, Chapter 15, Configuring Dynamic Interfaces*. When you create a profile for dynamic MLD interfaces, specify attributes for configuring all layers in the interface.

You use the MLD commands shown in Table 11 to configure a static MLD interface. You also use these commands to define the attributes for the MLD layer when you create a profile for dynamic MLD interfaces.

Table 11: Static MLD Commands

ipv6 mld	**ipv6 mld query-interval**
ipv6 mld access-group	**ipv6 mld query-max-response-time**
ipv6 mld access-source-group	**ipv6 mld robustness**
ipv6 mld explicit-tracking	**ipv6 mld static-include**
ipv6 mld group limit	**ipv6 mld static-exclude**
ipv6 mld immediate-leave	**ipv6 mld static-group**
ipv6 mld last-member-query-interval	**ipv6 mld version**
ipv6 mld querier-timeout	

The following sections describe the tasks associated with these and other **ipv6 mld** commands.

You can also use various MLD-specific RADIUS attributes in RADIUS Access-Accept messages as an alternative method of configuring certain values. See *JUNOSe Broadband Access Configuration Guide, Chapter 3, Configuring RADIUS Attributes* for additional information.

Enabling MLD on an Interface

You must start MLD on each interface that you want to use the protocol. You can configure MLD and PIM on the same interface. If you configure only MLD on an interface, the router determines that MLD owns that interface. If you configure MLD and PIM on an interface, the router determines that PIM owns the interface.

In an MLDv1 or MLDv2 network, the querier is the router with the lowest IPv6 address.

To start MLD, complete the following steps:

1. Enable MLD on the interface (MLDv2 is the default version).
2. (MLDv1) Specify the MLD version for the interface.

ipv6 mld

- Use to enable MLD on an interface and to set the MLD version to MLDv2. Use the **ipv6 mld version** command to specify a different MLD version.
- Example

  host1:boston(config-if)#**ipv6 mld**

- Use the **no** version to disable MLD on an interface.

ipv6 mld version

- Use to set the MLD version (1 or 2) for the interface.
- Example

  host1:boston(config-if)#**ipv6 mld version 2**

- Use the **no** version to set the version to the default, MLDv2.

Configuring MLD Settings for an Interface

When you start MLD on an interface, it operates with the default settings. You can, however, modify:

- The method that the router uses to remove hosts from multicast groups
- The time interval at which the querier sends multicast listener queries
- The time that a querier waits before sending a new query to hosts from which it receives multicast listener done messages
- The time that a non-querier waits for queries from the current querier before sending query messages to assume responsibility of querier
- The time that a host can take to reply to a query (maximum response time)
- The number of times that the router sends each MLD message from this interface

ipv6 mld immediate-leave

- Use to specify that, when the router receives a multicast listener done message from a host associated with this interface, the router immediately removes that host from the multicast group.

CAUTION: Issue this command only on MLD interfaces to which one MLD host is connected. If more than one MLD host is connected to a LAN through the same interface, and one host sends a done message, the router removes all hosts on the interface from the multicast group. The router loses contact with the hosts that should remain in the multicast group until they send join requests in response to the router's next general multicast listener query.

- Use the MLD-Immediate-Leave RADIUS attribute (VSA 26-100) in RADIUS Access-Accept messages as an alternative method of configuring this value. The RADIUS setting takes precedence over a CLI setting.
- Example

  host1:boston(config-if)#**ipv6 mld immediate-leave**

- Use the **no** version to restore the default behavior, in which the router removes a host from a multicast group if that host does not return a multicast listener report within a certain length of time after receiving a multicast listener query from the router.

ipv6 mld last-member-query-interval

- Use to specify the last-member-query-interval value, in the range 1–255 tenths of a second. When the router receives an MLDv1 leave message or an MLDv2 state change report, it sends out a query and expects a response within the time specified by this value.
- Using a lower value allows members to leave groups more quickly.
- Example

  host1:boston(config-if)#**ipv6 mld last-member-query-interval 90**

- Use the **no** version to restore the default, 10-tenths of a second (1 second).

ipv6 mld querier-timeout

- Use to set the time, in the range 1–400 seconds, that the interface waits for queries from the current querier before sending query messages to assume responsibility of querier.
- Example

  host1:boston(config-if)#**ipv6 mld querier-timeout 200**

- Use the **no** version to set the time to the default, twice the query interval.

ipv6 mld query-interval

- Use to specify how often, in the range 1–300 seconds, the interface sends group membership queries.
- Use the MLD-Query-Interval RADIUS attribute (VSA 26-98) in RADIUS Access-Accept messages as an alternative method of configuring this value. The RADIUS setting takes precedence over a CLI setting.
- Example

  host1:boston(config-if)#**ipv6 mld query-interval 100**

- Use the **no** version to set the polling interval to the default, 125 seconds.

ipv6 mld query-max-response-time

- Use to specify the period in tenths of a second during which the host is expected to respond to a group membership query. The possible period ranges are as follows:
  - IGMPv1 and IGMPv2: 1–255 tenths of a second
  - IGMPv3: 1–31744 tenths of a second
- MLDv1 and MLDv2 include this value in MLD query messages sent out on the interface.
- Using a lower value allows members to join and leave groups more quickly.

- Use the MLD-Query-Max-Resp-Time RADIUS attribute (VSA 26-99) in RADIUS Access-Accept messages as an alternative method of configuring this value. The RADIUS setting takes precedence over a CLI setting.
- Example

  host1:boston(config-if)#**ipv6 mld query-max-response-time 120**

- Use the **no** version to restore the default, 100 tenths of a second (10 seconds).

ipv6 mld robustness

- Use to specify the number of times that the router sends each MLD message from this interface.
- Use a higher value to ensure high reliability from MLD.
- Specify a number in the range 1–4.
- Example

  host1:boston(config-if)#**ipv6 mld robustness 2**

- Use the **no** version to restore the default, 3.

Specifying Multicast Groups

You can use a standard IPv6 access list to specify the multicast groups that a host can join.

ipv6 mld access-group

- Use to restrict hosts on this subnetwork to join only multicast groups that appear on the specified IPv6 access list.
- When configured, the access list is queried whenever the router receives an MLDv1 report requesting membership of a group and MLDv2 ChangeToInclude, IsInclude, ChangeToExclude, or IsExclude reports. The request is ignored if the access list query fails. The **ipv6 mld access-group** command uses IPv6 access lists, which allow both source and destination/group addresses to be specified. You must set the source address to "any."
- Use the MLD-Access-Name RADIUS attribute (VSA 26-74) in RADIUS Access-Accept messages as an alternative method of configuring this value.
- Example

  host1:boston(config-if)#**ipv6 mld access-group boston-list**

- Use the **no** version to dissociate the interface from an access list and to allow hosts on the interface to join any multicast group.

ipv6 mld access-source-group

- Use to restrict hosts on this subnetwork to membership in those (S,G) pairs (also known as "channels") permitted by the specified IPv6 access list.
- When configured, both source and group addresses query the associated access list whenever the router receives an MLDv2 report requesting membership of the (S,G) pairs (that is, the router receives an MLDv2 ChangeToInclude, IsInclude, or AllowNewSource group report). The request is ignored if the access list query fails. The **ipv6 mld access-source-group** command uses IPv6 access lists, which allow both source and destination group addresses to be specified.
- Use the MLD-Access-Src-Name RADIUS attribute (VSA 26-75) in RADIUS Access-Accept messages as an alternative method of configuring this value.
- Example

  host1:boston(config-if)#**ipv6 mld access-source-group dallas-list**

- Use the **no** version to remove any access list restriction.

Assigning a Multicast Group to an Interface

You can assign an interface to send and receive all traffic for a particular multicast group. This feature allows you to control the MLD traffic and to test the behavior of multicast protocols in the network.

ipv6 mld static-group

- Use to send and receive all traffic for a multicast group from a specific interface.
- The interface sets no timers for this group.
- Example

  host1:boston(config-if)#**ipv6 mld static-group ff0e::1**

- Use the **no** version to remove the group from the interface.

Configuring Group Outgoing Interface Mapping

You can configure an MLD protocol interface to use a different outgoing interface (OIF) for multicast-data-forwarding by applying an OIF map. When you configure an OIF map on an MLD protocol interface, the map is applied to all MLD membership requests that the interface receives. To configure OIF mapping on an interface, you first create the OIF map using the **ipv6 mld oif-map** command and then apply that map to an interface with the **ipv6 mld apply-oif-map** command.

To properly configure an interface used in the OIF map for multicast-data-forwarding capability, you must configure the interface version as passive with the **ipv6 mld version** command. You can either specify a passive interface as the OIF or specify the OIF as *self* (to use the MLD protocol interface as the OIF) in the **ipv6 mld oif-map** command.

ipv6 mld apply-oif-map

- Use to apply the specified outgoing interface (OIF) map to the current interface.
- Example

  host1(config-subif)#**ipv6 mld apply-oif-map OIFMAP**

- Use the **no** version to remove the outgoing interface map association from the interface.

ipv6 mld oif-map

- Use to create an OIF map.
- Use the MLD-OIF-Map-Name RADIUS attribute (VSA 26-76) in RADIUS Access-Accept messages as an alternative method of configuring this value.
- Example

host1(config)#**ipv6 mld oif-map OIFMAP atm 3/0.1 ff0e::1:1/128 2001::1:1/128**
host1(config)#**ipv6 mld oif-map OIFMAP atm 3/0.2 ff0e::1:1/128 2001::1:1/128**
host1(config)#**ipv6 mld oif-map OIFMAP atm 3/0.3 ff0e::1:1/128**
host1(config)#**ipv6 mld oif-map OIFMAP atm 3/0.4 ff0e::1:0/112 2001::1:1/128**
host1(config)#**ipv6 mld oif-map OIFMAP atm 3/0.5 ff0e::1:0/112 2001::1:1/128**
host1(config)#**ipv6 mld oif-map OIFMAP self ::/0 2001::1:0/112**

- Use the **no** version to remove an outgoing interface map attribute.

ipv6 mld version

- Use to set the MLD version (1 or 2) for the interface or specify a passive interface with only multicast-data-forwarding capability (passive).
- Use the MLD-Version RADIUS attribute (VSA 26-77) in RADIUS Access-Accept messages as an alternative method of configuring this value.
- Example

  host1:dallas(config-if)#**ipv6 mld version passive**

- Use the **no** version to set the version to the default, mldv2.

Configuring SSM Mapping

SSM mapping enables the router to determine one or more source addresses for group G. The mapping effectively translates an MLDv1 multicast listener report to an MLDv2 report, enabling the router to continue as if it had initially received an MLDv2 report. After the router is joined to these groups, it sends out PIM join messages and continues to enable joining from these groups, as long as it continues to receive MLDv1 membership reports and no change occurs to the SSM mapping for the group.

When you statically configure SSM mapping, the router can discover source addresses from a statically configured table.

The following applies when you configure SSM mapping:

- When enabled, and either you have not configured a static SSM map or the router cannot find any matching access lists, the router continues to accept (*,G) groups. The PIM SSM range must deny any unacceptable SSM group addresses.
- When you issue the **no ipv6 mld ssm-map enable** command, the router removes all SSM map (S,G) states and establishes a (*,G) state.
- You can enter multiple **ssm-map static** commands for different access lists. Also, you can enter multiple **ssm-map static** commands for the same access list, as long as the access list uses different source addresses.
- SSM maps do not process statically configured groups.

ipv6 mld ssm-map enable

- Use to enable SSM mapping on the router. SSM mapping statically assigns sources to MLDv1 groups. You must use SSM mapping for MLDv1 hosts to interoperate with PIM SSM. SSM mapping allows the router use a statically configured list to translate (*,G) memberships to (S,G) memberships.
- Example

  host1:boston(config)#**ipv6 mld ssm-map enable**

- Use the **no** version to disable the SSM map.

ipv6 mld ssm-map static

- Use to specify an access list and source address for use in SSM mapping. SSM mapping statically assigns sources to MLDv1 groups. You must use SSM mapping for MLDv1 hosts to interoperate with PIM SSM. SSM mapping allows the router to use a statically configured list to translate (*,G) memberships to (S,G) memberships.
- The **ipv6 mld ssm-map static** command uses IPv6 access lists, which allow both source and destination/group addresses to be specified. You must set the source address to "any."
- Example

  host1:boston(config)#**ipv6 mld ssm-map static boston-list 2001::1**

- Use the **no** version to remove the SSM map association.

Limiting the Number of Accepted MLD Groups

By default, there is no limit on the number of MLD groups that an MLD interface can accept. However, you can manage multicast traffic on the router by restricting the number of MLD groups accepted by:

- A specific port on an I/O module
- A specific MLD interface

If you set limits for both a port and interfaces on that port, the router uses the lower of the two limits when determining how many MLD groups an interface can accept. For example, if you set a limit of 10 groups for the port and 15 groups for each interface, the router allows only 10 groups to be accepted among the interfaces.

However, if you set a limit for a port and that limit is lower than the number of groups currently accepted by the interfaces on that port, the router does not dissociate the groups from the interfaces. The router enforces the new limit on the port when the number of groups associated with the interfaces falls to that limit. For example, if the interfaces on the port have accepted a total of 15 groups, and you set a limit of 10 groups on the port, the router does not disconnect any of the groups and does not allow the interfaces to accept any more groups. Over time, some groups leave the interfaces and, eventually, a maximum of ten groups remains connected.

ipv6 mld group limit

- Use to limit the number of MLD groups that an interface can accept.
- Example

  host1:boston(config-if)#**ipv6 mld group limit 5**

- Use the **no** version to restore the default situation, in which there is no limit on the number of MLD groups that an interface can accept.

multicast group port limit

- Use to limit the number of MLD groups that a port can accept.
- Specify the identifier for the port in *slot/port* format (ERX routers) or in *slot/adapter/port* format (E120 and E320 routers) and the maximum number of MLD groups that interfaces can accept.
  - *slot*—Number of the chassis slot in the range 0–6 (ERX-7xx models), 0–13 (ERX-14xx models), 0–5 (E120 router), or 0–16 (E320 router)
  - *adapter*—Number of the bay in which the I/O adapter (IOA) resides. This identifier applies to the E120 and E320 routers only. In the software, adapter 0 identifies the right IOA bay (E120 router) and the upper IOA bay (E320 router); adapter 1 identifies the left IOA bay (E120 router) and the lower IOA bay (E320 router).
  - *port*—Port number on the I/O module or IOA
- Example 1—ERX models

  host1(config)#**multicast group port 3/0 limit 5**

- Example 2—E120 and E320 routers

  host1(config)#**multicast group port 3/1/0 limit 5**

- Use the **no** version to restore the default situation, in which there is no limit on the number of MLD groups that a port can accept.

Including and Excluding Traffic

MLDv2 extends MLDv1 functionality with the ability to include or exclude specific multicast traffic sources. That is, with MLDv2, hosts signal (S,G) pairs that they want to include or exclude.

For hosts that cannot signal group membership dynamically, you can use the **ipv6 mld static-include** or **ipv6 mld static-exclude** command to statically include or exclude multicast traffic, respectively.

MLDv2 is the industry-designated standard protocol for hosts to signal channel subscriptions in SSM. For additional information about SSM, see *PIM Source-Specific Multicast* on page 103.

ipv6 mld static-exclude

- Use to statically exclude the MLD (S,G) membership for a host that is not capable of dynamically signaling group membership.
- Example

  host1:boston(config-if)#**ipv6 mld static-exclude 2001::1 ff0e::1**

- Use the **no** version to remove the static designation.

ipv6 mld static-include

- Use to statically include the MLD (S,G) membership for a host that is not capable of dynamically signaling group membership.
- Example

  host1:boston(config-if)#**ipv6 mld static-include 2001::1 ff0e::1**

- Use the **no** version to remove the static designation.

Configuring Explicit Host Tracking

Explicit host tracking enables the router to explicitly track each individual host that is joined to a group or channel on a particular multi-access network.

Explicit host tracking provides the following benefits:

- Minimal leave latency when a host leaves a multicast group or channel. When the router receives a leave message for a group or channel on an interface, it accesses a list of hosts and immediately stops forwarding traffic if the sender is the last host to request traffic for that group or channel. The leave latency is bound only by the packet transmission latencies in the multi-access network and the processing time in the router.
- Ability to change channels quickly in networks where bandwidth is constrained between a multicast-enabled router and hosts.
- Ability to determine what multicast hosts are joined to particular multicast groups or channels; this is useful for accounting purposes.
- Reduction of control message traffic on the network because, when it receives a leave message, the router no longer needs to send out MLD queries to verify membership. As a result, interested hosts also do not need to respond to these queries with reports.
- Tracking based on MLD reports for hosts in both include and exclude modes for every multicast group or channel on an interface.

When the router is configured for explicit host tracking and starts performing immediate leave using the host information collected, every leave message received for a group or channel is treated as follows:

- The router checks the number of hosts that receive traffic from this group or channel.
- If the host sending the leave message is the only host, it performs immediate leave for that group or channel on that interface. The router removes the interface from the multicast group or channel immediately, without sending out a group or group-source specific query and waiting for the last member query interval.
- If the host sending the leave message is not the only host receiving traffic for that group or channel, the router removes the host from the list of hosts on that interface, but keeps the interface in the outgoing interface list for the multicast group or channel. No group or group-source specific queries are sent.

You can enable MLD explicit host tracking on an interface only if MLD V1 or V2 has been previously enabled on the interface. Explicit host tracking is not enabled by default when you enable MLD on the interface. Explicit host tracking cannot be configured on passive MLD interfaces.

When you enable explicit host tracking on an interface that has a membership state, the router does not immediately start performing immediate leave. For a maximum of group membership interval seconds, the router only performs host tracking. Any leave messages that the router receives during this period receive normal leave processing. Any leave messages received after this interval has elapsed receive immediate leave processing, when appropriate.

When explicit host tracking is enabled on an MLD V2 interface, even if a group has to downgrade to MLD V1 due to the presence of an MLD V1 host, explicit host tracking continues for that group. To avoid this, you can use the **disable-if-mld-v1-detected** keyword. If you select this option, the router turns off explicit host tracking for the group when MLD V1 host reports are received for the group on that interface. This option does not have any significance on an interface configured for MLD V1 and is ignored if selected.

If you execute the command on an interface that was previously enabled for immediate-leave, the configuration is accepted, immediate-leave is turned off and an appropriate warning message logged. Any attempt to configure immediate-leave on an interface that has explicit host tracking enabled is rejected and an error message logged.

The following example enables MLD V2 explicit host tracking on interface 3/0.101 with the default configuration where the router continues to perform explicit host tracking for MLD V1 groups. To override this default configuration, use the **ip mld explicit-tracking disable-if-mld-v1-detected** command.

```
interface 3/0.101
ip mld version 2
ip mld explicit-tracking
end
```

ipv6 mld explicit-tracking

- Use to set explicit host tracking for MLD interfaces.
- To disable explicit host tracking if MLD V1 hosts are detected, use the **disable-if-mld-detected** keyword.
- Example

  host1(config)#**ipv6 mld explicit-tracking**

- Use the **no** version to disable explicit host tracking on the interface. Use the **no** version with the **disable-if-mld-detected** keyword to revert to the default explicit host tracking behavior.

Disabling and Removing MLD

You can disable and reenable MLD on the VR. You can also remove MLD from the VR and re-create it on the VR.

mld disable

- Use to disable MLD on a VR.
- Example

  host1(config)#**virtual-router boston**
  host1:boston(config)#**router mld**
  host1:boston(config-router)#**mld disable**

- Use the **no** version to enable MLD on a VR.

router mld
ipv6 router mld

- Use to create and enable MLD on a VR or to access MLD Router Configuration mode.
- Example 1

  host1(config)#**virtual-router boston**
  host1:boston(config)#**router mld**

- Example 2

  host1(config)#**virtual-router boston**
  host1:boston(config)#**ipv6 router mld**

- Use the **no** version to delete MLD and MLD proxy from the VR.

Monitoring MLD

You can establish a reference point for MLD statistics by setting the statistics counters to zero.

To display MLD parameters, use the **show** commands described in this section.

NOTE: The E120 router and E320 router output for **monitor** and **show** commands is identical to output from other E-series routers, except that the E120 and E320 router output also includes information about the adapter identifier in the interface specifier (*slot/adapter/port*).

baseline ipv6 mld

- Use to set the counters for MLD statistics to zero.
- Example

  (host1)#**baseline ipv6 mld**

- There is no **no** version.

show ipv6 mld

- Use to display MLD information for a VR.
- Field descriptions
  - Administrative state—Status of MLD in the software: enabled or disabled
  - Operational state—Status of MLD on the VR: enabled or disabled
  - total interfaces—Number of interfaces on which you started MLD
  - enabled—Number of interfaces on which MLD is enabled
  - disabled—Number of interfaces on which MLD is disabled
  - learned groups—Number of multicast groups that the VR has discovered
  - MLD Statistics Rcvd—Statistics for MLD messages received
    - total—Number of MLD messages received
    - checksum errors—Number of MLD messages received with checksum errors
    - unknown types—Number of messages received that are not multicast listener queries, multicast listener reports, or multicast listener done messages
    - discards—Number of multicast listener discards
    - queries—Number of multicast listener queries
    - reports—Number of multicast listener reports
    - leaves—Number of done messages
  - MLD Statistics Sent—Number of multicast listener queries sent
- Example

```
host1:boston#show ipv6 mld
Routing Process MLD, Administrative state enabled, Operational state enabled
    4 total interfaces, 4 enabled, 0 disabled
    2 learned groups
MLD Statistics:
    Rcvd: 3 total, 0 checksum errors, 0 unknown types, 0 discards
          0 queries, 3 reports, 0 leaves
    Sent: 5 total
```

show ipv6 mld groups

- Use to display statically joined and directly connected groups learned through MLD.
- Field descriptions
  - Grp Address—Address of the multicast group
  - Interface—Interface that discovered the multicast group
  - oif-map—Name of the OIF map and the mapped OIF interface, if a group or source has been mapped to an OIF
  - State—MLD version of the group
  - Reporter—Link-local address of the host reporting the multicast group
  - ExpTim—Remaining time, in seconds, at which the router stops polling for more members of this group
  - oldHTo—Remaining time at which the router stops polling for more MLDv1 members of a group. If this value is 0, the interface has received no MLDv1 reports for the group.
  - Included Sources—Sources included in the multicast group
  - Excluded Sources—Sources excluded from the multicast group
  - Counts—Number of source-group mappings by version and state
- Example 1—Without OIF mapping

```
host1:boston#show ipv6 mld groups
    Grp Address           Interface          State      Reporter     ExpTim oldHTo
------------------- --------------------- -------- --------------- ------ ------
ff0e::1             ATM2/0.15             Version2 fe80::90:1a02:1 54     0
                                                   640:91d
ff0e::4:1           ATM2/0.15             Version2 fe80::90:1a02:1 54     0
                                                   640:91d
  Included Sources:
    51::1                                                          54
    51::2                                                          54

Counts: 2 version-2, 0 version-1, 0 check state, 0 disabled
       (2 total)
        0 excluded
        Source-groups: 2 included, 0 excluded
```

- Example 2—With OIF mapping

```
host1:boston#show ipv6 mld groups
    Grp Address           Interface          State      Reporter     ExpTim oldHTo
------------------- --------------------- -------- --------------- ------ ------
ff3e::1             ATM5/0.12             Version2 fe80::f7:0:91a: 377    0
                                                   0
                    oif-map OIFMAP ATM5/0.
                    121
ff3e::1             ATM5/0.13             Version2 fe80::f7:0:a1a: 369    0
                                                   0
                    oif-map OIFMAP ATM5/0.
                    121
ff3e::2             ATM5/0.12             Version2 fe80::f7:0:91a: 370    0
                                                   0
```

```
  Included Sources:
    10::2           oif-map OIFMAP self                          370
    10::10          oif-map OIFMAP ATM5/0.                       370
                    120
    10::11          oif-map OIFMAP ATM5/0.                       370
                    121
ff3e::2             ATM5/0.13             Version2 fe80::f7:0:a1a: 373    0
                                                   0
  Included Sources:
    10::2           oif-map OIFMAP self                          373
    10::10          oif-map OIFMAP ATM5/0.                       373
                    120
    10::11          oif-map OIFMAP ATM5/0.                       373
                    121

Counts: 4 version-2, 0 version-1, 0 check state, 0 disabled
       (4 total)
        0 excluded
        Source-groups: 6 included, 0 excluded
```

show ipv6 mld interface

- Use to display MLD information for interfaces on which you enabled MLD.
- Specify the **brief** keyword to see a summary of the information.
- Specify the **count** keyword to see the number of MLD interfaces.
- Field descriptions
  - Interface—Type of interface and interface specifier. For details about interface types and specifiers, see *JUNOSe Command Reference Guide, About This Guide*.
  - address—IPv6 link-local address of the interface
  - Administrative state—Status of the interface in the software: enabled or disabled
  - Operational state—Physical status of the interface: enabled or disabled
  - Version—MLD version
  - State—Function of the interface: querier or nonquerier
  - Query Interval—Time interval at which this interface sends query messages
  - Other querier present interval—Time that the interface waits before declaring itself as the querier
  - Maximum response time—Time interval during which this interface expects a host to respond
  - Last member query interval—Time that this interface waits before sending a new query to a host that sends a group leave message
  - Robustness—Number of times this interface sends MLD messages
  - Information about IPv6 access lists configured with the **ipv6 mld access-group** command
    - Inbound access group—Access list specified
    - No inbound access group—No access list specified

- Information about IPv6 access lists configured with the **ipv6 mld access-source-group** command
  - Inbound access source-group—Access list specified
  - No inbound access source-group—No access list specified
- Information about OIF maps configured with the **ipv6 mld apply-oif-map** command
  - Inbound apply-oif-map—Map name specified
  - No inbound apply-oif-map—No map name specified
- Immediate Leave—Setting of the **ipv6 mld immediate-leave** command: enabled or disabled
- Explicit Host Tracking—Setting of the **ipv6 mld explicit-tracking** command: enabled or disabled
- Max-Group limit—Number of MLD groups that the interface can accept, as configured with the **ipv6 mld group limit** command
- Group Count—Number of MLD groups that the interface has accepted
- IOA packet replication—Hardware multicast packet replication interface to which egress multicast packets on this interface are redirected
- Interface statistics Rcvd—Information about MLD messages received on this interface
  - reports—Number of group multicast listener reports received
  - leaves—Number of group multicast listener done messages received
  - wrong version queries—Number of multicast listener queries received from devices running a different version of MLD
- Interface statistics Sent—Number of MLD messages this interface has sent
- Interface statistics Groups learned—Number of groups this interface has discovered
- Counts—Total number of MLD interfaces

- Example

```
host1:boston#show ipv6 mld interface

Interface ATM5/0.1 address fe80::f7:0:321a:0
Administrative state enabled, Operational state enabled
Interface parameters:
    Version 1
    State Querier
    Query Interval 125 secs, 123 secs before the next query
    Other querier present interval 250 secs
    Maximum response time 100 (in 10ths of a second)
    Last member query interval 10 (in 10ths of a second)
    Robustness 3
    No inbound access group
    No inbound access source-group
    No inbound apply-oif-map
    Immediate Leave: disabled
    Explicit Host Tracking: enabled
    Max-Group limit: No Limit
    Group Count: 0
```

```
Interface statistics:
    Rcvd: 0 reports, 0 leaves, 0 wrong version queries
    Sent: 14 queries
    Groups learned: 0

Counts: 0 down, 0 init state, 1 querier, 0 non-querier, 1 Total
```

- Example 2

```
host1#show ipv6 mld interface gigabitEthernet 3/0.0
Interface GigabitEthernet3/0.0 address 10.1.1.1/255.255.255.0
Administrative state enabled, Operational state enabled
Interface parameters:
    Version 1
    State Querier
    Query Interval 125 secs, 123 secs before the next query
    Other querier present interval 250 secs
    Maximum response time 100 (in 10ths of a second)
    Last member query interval 10 (in 10ths of a second)
    Robustness 3
    No inbound access group
    No inbound access source-group
    No inbound apply-oif-map
    Immediate Leave: disabled
    Explicit Host Tracking: enabled
    Max-Group limit: No Limit
    Group Count: 0
    IOA packet replication gigabitEthernet 3/8.1
Interface statistics:
    Rcvd: 0 reports, 0 leaves, 0 wrong version queries
    Sent: 14 queries
    Groups learned: 0

Counts: 0 down, 0 init state, 1 querier, 0 non-querier, 1 Total
```

show ipv6 mld interface brief

- Use to display a summary of MLD information for interfaces on which you enabled MLD.
- Field descriptions
  - Interface—Type of interface and interface specifier. For details about interface types and specifiers, see *JUNOSe Command Reference Guide, About This Guide*.
  - Intf Address—IPv6 link-local address of the interface
  - Ver—MLD version
  - State—Function of the interface: querier or nonquerier
  - Querier—IPv6 address of the querier on the network to which this interface connects
  - QTime—Remaining time interval at which this interface sends query messages
  - QPTime—Remaining time that the interface waits before declaring itself as the querier

- Example

```
host1:boston#show ipv6 mld interface brief

   Interface     Intf Address   Ver    State        Querier     QTime  QPTime
--------------- --------------- --- ----------- --------------- ------ ------
ATM5/0.1        fe80::f7:0:231a 1   Querier     fe80::f7:0:231a 1      0
                :0                              :0
ATM5/0.200      fe80::f7:0:231a 2   Querier     fe80::f7:0:231a 20     0
                :0                              :0

Counts: 0 down, 0 init state, 2 querier, 0 non-querier, 2 Total
```

show ipv6 mld mapped-oif

- Use to display the current mappings to all mapped outgoing interfaces or to the specified mapped outgoing interface.
- Field descriptions
  - OIF—Outgoing interface used in an OIF map
  - Oper—Operation status of the outgoing interface
  - Group Address—Multicast group IP address associated with the OIF
  - Source Address—Source IP address associated with the OIF
  - Join I/F—MLD protocol interface associated with the OIF
  - Map Name—Name of the map associated to the OIF
  - Counts—Number of source-group mappings to OIFs
- Example

```
host1#show ipv6 mld mapped-oif
      OIF       Oper  Group Address  Source Address     Join I/F      Map Name
--------------- ---- --------------- --------------- --------------- ----------
ATM5/0.120      Up   ff3e::2         10::10          ATM5/0.12       OIFMAP
                                                     ATM5/0.13       OIFMAP
ATM5/0.121      Up   ff3e::1         *               ATM5/0.12       OIFMAP
                                                     ATM5/0.13       OIFMAP
                     ff3e::2         10::11          ATM5/0.12       OIFMAP
                                                     ATM5/0.13       OIFMAP

Counts: 3 source-group mappings
```

show ipv6 mld oif-map

- Use to display all outgoing interface (OIF) maps or the OIF map for the specified map name.
- Field descriptions
  - Map Name—Name of the map associated to the show output
  - Group Prefix—Multicast group IPv6 prefix
  - Source Prefix—Source IPv6 prefix
  - OIF—Outgoing interface associated with the group and source prefix
- Example

```
host1#show ipv6 mld oif-map
     Map Name         Group Prefix      Source Prefix           OIF
------------------ ------------------ ------------------ ------------------
OIFMAP             ff3e::/112         ::/0               ATM5/0.121
                   ff3e::/112         10::2/128          self
                   ff3e::/112         10::10/128         ATM5/0.120
                   ff3e::3/128        ::/0               ATM5/0.130
                   ff3e::4/128        ::/0               ATM5/0.130
```

show ipv6 mld membership

- Use to display MLD membership information for multicast groups and (S, G) channels.
- Specify the **brief** keyword to see a summary of the information.
- Specify the **tracked** keyword to see interface information only for interfaces where explicit host tracking is enabled.
- Field descriptions
  - Group—Multicast group or (S, G) channel
  - Source—(S, G) entries that are forwarding traffic
  - Reporter—Hosts that requested including sources or that have not requested excluding sources. If listed under a group, host that sent exclude reports for the group. If listed under a source, host that requested traffic from this source for the group. For any (S, G), if listed under a source, indicates hosts interested in the traffic for this (S, G).

- ExpTim—Expiration time
- Flags
  - M—Uses Oifmap
  - S—SSM mapped
  - T—Tracked
  - 1, 2—MLD version that the group is in
- Interface—Type of interface and interface specifier. For details about interface types and specifiers, see *JUNOSe Command Reference Guide, About This Guide*.

- Example

```
host1:boston#show ipv6 mld membership
Flags: M - Uses Oifmap  S- SSM mapped  T - tracked
       1,2 - The version of MLD the group is in
Reporter:
    <ip-address> - last reporter if the group is not explicitly tracked
    <n>/<m> - <n> reporters include mode, <m> reporters in exclude

Group      Source      Reporter                 ExpTim    Flags     Interface
---------  --------    ----------------------   ------    ------    ---------------
ff0e::40      *        fe80::90:1a02:1640:91d   02:41     2S        FastEthernet2/1

ff0e::50               1/2                      02:56    3MT       FastEthernet2/2
                       fe80::90:1a02:1640:911   02:30
                       fe80::90:1a02:1640:912   02:48
           20::11
                       fe80::90:1a02:1640:913   02:56
           20::12
                       fe80::90:1a02:1640:911   02:30
           20::13
                       fe80::90:1a02:1640:911   02:30
                       fe80::90:1a02:1640:912   02:48
                       fe80::90:1a02:1640:913   02:56
ff0e::60               fe80::90:1a02:1640:901   01:56     3         FastEthernet2/3
           10::10                               02:45
           10::11                               02:35
           10::12                               02:15
           10::14                               stop
ff0e::70               fe80::90:1a02:1640:91    stop      3         FastEthernet2/4
           40::10                               01:10
           40::11                               01.24
ff0e::80               2/0                      stop      3T        FastEthernet2/5
           50::10
                       fe80::90:1a02:1650:910   02:48
           50::11
                       fe80::90:1a02:1650:920   02:56
                       fe80::90:1a02:1650:910   02:48
           50::12      fe80::90:1a02:1650:920   02:56
ff0e::90               0/3                      02:56     2T        FastEthernet2/6

              *
                       fe80::90:1a02:1660:910   02:48
                       fe80::90:1a02:1660:920   02:56
                       fe80::90:1a02:1660:930   02:48
```

show ipv6 mld oif-mapping

- Use to display the mapped OIF to be assigned to a given map-name, group address, and source address.
- Field descriptions
  - OIF-MAP Name—Name of the map requested
  - Group Address—Multicast group IP address requested
  - Source Address—Source IP address requested
  - Mapped OIF—Join interface associated with the OIF map
- Example

```
host1#show ipv6 mld oif-mapping OIFMAP ff3e::1 10::10
    OIF Mapping
    OIF-MAP Name   : OIFMAP
    Group Address  : ff3e::1
    Source Address : 10::10
    Mapped OIF     : ATM5/0.120
```

show ipv6 mld ssm-mapping

- Use to display the SSM mapping state and the source list mapping associated with a multicast group address.
- Field descriptions
  - SSM Mapping—Status of SSM mapping on the interface (enabled or disabled)
  - Group Address—Multicast group address requested
  - Source List—List of sources mapped to the multicast group address
- Example

```
host1:boston#show ipv6 mld ssm-mapping ff3e::1
    SSM Mapping   : Enabled
    Group Address : ff3e::1
    Source List   : 2001::1
                  : 2001::2
```

show multicast group limit

- Use to display the number of MLD groups that ports have accepted and, if configured, the maximum number of groups that ports can accept.
- A value of –1 indicates that no port group limit is configured.
- Only ports that have accepted MLD groups and ports for which you have configured a limit for the number of MLD groups appear in this display.

- Field descriptions
  - Port—Identifier of the port in *slot*/*port* format
    - *slot*—Number of the chassis slot in the range 0–6 (ERX-7xx models) or 0–13 (ERX-14xx models)
    - *port*—Port number on the I/O module
  - limit—Maximum number of MLD groups that the port can accept. A value of –1 indicates that no limit has been specified.
  - count—Actual number of MLD groups the port has accepted
- Example

```
host1:boston#show multicast group limit
Port      limit count
--------- ----- -----
2/0           5     0
2/1          -1     1
```

MLD Proxy Overview

MLD proxy enables the router to issue MLD host messages on behalf of hosts that the router discovered through standard MLD interfaces. The router acts as a *proxy* for its hosts. The E-series router supports MLD proxy versions 1 and 2.

Figure 18 shows a router in an MLD proxy configuration. You enable MLD proxy on one interface, which connects to a router closer to the root of the tree. This interface is the *upstream interface*. The attached upstream router on the upstream interface should be running MLD.

You enable MLD on the interfaces that connect the router to its hosts that are farther away from the root of the tree. These interfaces are known as *downstream interfaces*.

Figure 18: Upstream and Downstream Interfaces

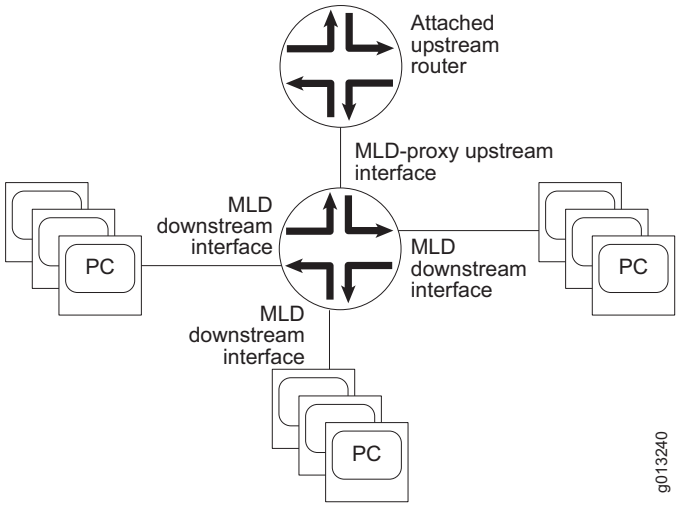

As described in *Overview* on page 194, hosts interact with the router through the exchange of MLD messages. Similarly, when you configure MLD proxy, the router interacts with the router on its upstream interface through the exchange of MLD messages. However, when acting as the proxy, the router performs the host portion of the MLD task on the upstream interface as follows:

- When queried, sends multicast listener reports to the group.
- When one of its hosts joins a multicast address group to which none of its other hosts belong, sends unsolicited multicast listener reports to that group.
- When the last of its hosts in a particular multicast group leaves, the group sends either an unsolicited multicast listener done report to the all-routers address (FF02::2) for MLDv1 or an MLDv2 multicast listener report to the all-MLDv2 routers address (FF02::16).

Configuring MLD Proxy

To configure a downstream interface, enable MLD on that interface. To configure MLD proxy on the router, complete the following tasks:

1. Enable IPv6 multicast.
2. Identify the interface that you want to act as the upstream interface.
3. Enable MLD proxy on that interface.
4. (Optional) Specify how often the router should send unsolicited reports to routers on the upstream interface.

ipv6 mld-proxy

- Use to enable MLD proxy on an interface.
- The interface for which you enable MLD proxy is the upstream interface.

NOTE: You can enable only one upstream interface.

- You can specify either MLD proxy version 1 or 2. The default is version 2.
- Example

  host1(config-if)#**ipv6 mld-proxy**

- Use the **no** version to disable MLD proxy on an interface.

ipv6 mld-proxy unsolicited-report-interval

- Use to specify, in tenths of a second, how often the upstream interface should transmit unsolicited reports.

NOTE: Issue this command only on the upstream interface. Otherwise, this command has no effect.

- Example

  host1(config-if)#**ipv6 mld-proxy unsolicited-report-interval 600**

- Use the **no** version to transmit unsolicited reports using the default value, 100-tenths of a second (10 seconds).

ipv6 mld-proxy version

- Use to set the MLD proxy version for the interface.
- Example

  host1(config-if)#**ipv6 mld-proxy version 1**

- Use the **no** version to set the version to its default value, MLDv2.

Setting the MLD Proxy Baseline

You can set the counters for the numbers of queries received and reports sent on the upstream interface to zero. This feature allows you to establish a reference point for MLD proxy statistics.

baseline ipv6 mld-proxy interface

- Use to set the counters for the numbers of queries received and reports sent on the upstream interface to zero.

NOTE: Issue this command only on the upstream interface. Otherwise, this command will have no effect.

- Example

  (host1)#**baseline ipv6 mld-proxy interface**

- There is no **no** version.

Monitoring MLD Proxy

To display MLD proxy parameters, use the following **show** commands.

show ipv6 mld-proxy

- Use to display MLD proxy parameters for a VR.
- Field descriptions
  - Routing Process—MLD proxy protocol
  - Administrative state—State of MLD proxy in the software
  - Operational state—Operational state of MLD proxy: enabled or disabled
  - total interfaces—Number of MLD proxy interfaces on the VR; currently only one upstream interface per VR
  - state—Operational state of the MLD proxy interfaces: enabled or disabled
  - multicast group—Number of multicast groups associated with MLD proxy interfaces
- Example

```
host1#show ipv6 mld-proxy
Routing Process MLD Proxy, Administrative state enabled, Operational state
enabled
    total 1 upstream interface, state enabled
    1 multicast group
```

show ipv6 mld-proxy groups

- Use to display information about multicast groups that MLD proxy reported.
- Field descriptions
  - Grp Address—Address of the multicast group
  - Interface—Type and identifier of the upstream interface associated with the multicast group
  - Grp Mode
    - Blank—No sources included or excluded for this group
    - Include—Sources included for this group
    - Exclude—Sources excluded for this group
  - Count—Total number of multicast groups associated with this interface
- Example 1

```
host1#show ipv6 mld-proxy groups

Grp Address     Interface       Grp Mode
--------------- --------------- --------------
ff0e::1         ATM5/1.200
ff0e::2         ATM5/1.200
ff0e::3         ATM5/1.200      Include(1):
  2001::1
```

```
ff0e::4         ATM5/1.200
ff0e::5         ATM5/1.200      Exclude(1):
  2001::2

Counts: 3 <*,G>, 1 Exclude (1 sources), 1 Include (1 sources)
       (5 total)
```

- Example 2

```
host1#show ipv6 mld-proxy groups ff0e::1
Grp Address     Interface       Grp Mode
--------------- --------------- --------------
ff0e::1         ATM5/1.200

Counts: 1 <*,G>
       (1 total)
```

- Example 3

```
host1#show ipv6 mld-proxy groups count
Counts: 3 <*,G>, 1 Exclude (1 sources), 1 Include (1 sources)
       (5 total)
```

show ipv6 mld-proxy interface

- Use to display information about the interface on which you configured MLD proxy.
- To view information about a particular interface, enter an interface type and specifier, such as atm 3/0. For details about interface types and specifiers, see *JUNOSe Command Reference Guide, About This Guide*.
- Specify the **brief** option to display a summary rather than a detailed description.
- Field descriptions
  - Interface—Type of upstream interface. For details about interface types, see *JUNOSe Command Reference Guide, About This Guide*.
  - Address—Address of upstream interface
  - Administrative state—State of upstream interface in the software: enabled or disabled
  - Operational state—Physical state of upstream interface: enabled or disabled
  - Version—MLD version on this interface
  - State—Presence of MLDv1 routers on the same subnet as this upstream interface
  - Unsolicited report interval—Time interval at which this upstream interface sends unsolicited group membership report
  - multicast group—Number of multicast groups associated with this upstream interface

- Interface statistics Rcvd—Statistics for messages received on this interface
  - v1 queries—Number of MLDv1 multicast listener queries received
  - v1 report—Number of MLDv1 multicast listener reports received
  - v2 queries—Number of MLDv2 multicast listener queries received
  - v2 report—Number of MLDv2 multicast listener reports received
- Interface statistics Sent—Statistics for messages sent from this interface
  - v1 reports—Number of MLDv1 multicast listener reports sent
  - v1 leaves—Number of multicast listener done messages sent
  - v2 reports—Number of MLDv2 multicast listener reports sent

- Example 1

```
host1#show ipv6 mld-proxy interface

Interface ATM5/1.200 address fe80::f7:0:231a:0
Administrative state enabled, Operational state enabled
Interface parameters:
    Version 2
    State No v1 Router Present
    Unsolicited report interval 100 (in 10ths of a second)
    5 multicast groups
Interface statistics:
    Rcvd: 0 v1 query, 0 v1 report, 25 v2 queries, 0 v2 report
    Sent: 0 v1 report, 0 v1 leave, 35 v2 reports
```

- Example 2

```
host1#show ipv6 mld-proxy interface brief

Interface       Intf Address       Ver State                UnSlTime
--------------- ------------------ --- -------------------- --------
ATM5/1.200      fe80::f7:0:231a:0    2 No v1 Router Present      100
```