

2 IP Multicast Routing

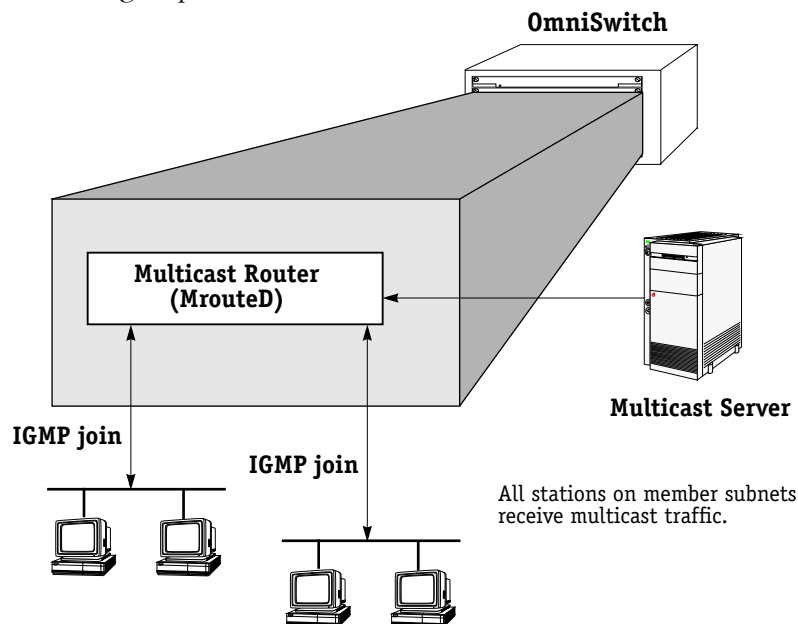
Introduction

This chapter describes how to install and use multicast routing on your switch. Typically, multicast routing is used in conjunction with IP Multicast Switching (IPMS). **Note that routing performance is decreased if you do not running IPMS with multicast routing.** For a brief overview of IP multicast and more details about IPMS, see Chapter 1, “IP Multicast Switching.”

Multicast uses Class D IP addresses in the range 224.0.0.0 to 239.255.255.255. Addresses in the range 239.0.0.0 to 239.255.255.255 are reserved for boundaries, which are used to prevent multicast traffic from being forwarded on a VLAN group or network (see *Modifying the Boundary List* on page 2-14). The multicast group address is indicated in the Destination Address field of the IP header. A multicast group is *not* a mobile or non-mobile group as defined in your switch user manual. In this chapter, mobile or non-mobile groups are referred to as *VLAN groups* to distinguish them from multicast groups.

IP multicast routing is a way of controlling multicast traffic across networks. The multicast router discovers which networks want to receive multicast traffic by sending out Internet Group Management Protocol (IGMP) queries and receiving IGMP reports from attached networks. The IGMP reports signal that users want to join or leave a multicast group. (Version 2 of IGMP is supported on the OmniSwitch, OmniStack, and OmniSwitch/Router.) If there is more than one multicast router in the network, the router with the lowest IP address is elected the querier router, which is responsible for querying the subnetwork for group members. (See RFC 2236 for more information about IGMP Version 2.)

The illustration given here shows two networks with end stations that send IGMP messages to a multicast router attached to a multicast server. All end stations attached to the networks receive multicast traffic, regardless of whether or not the particular end station requested to join the multicast group.

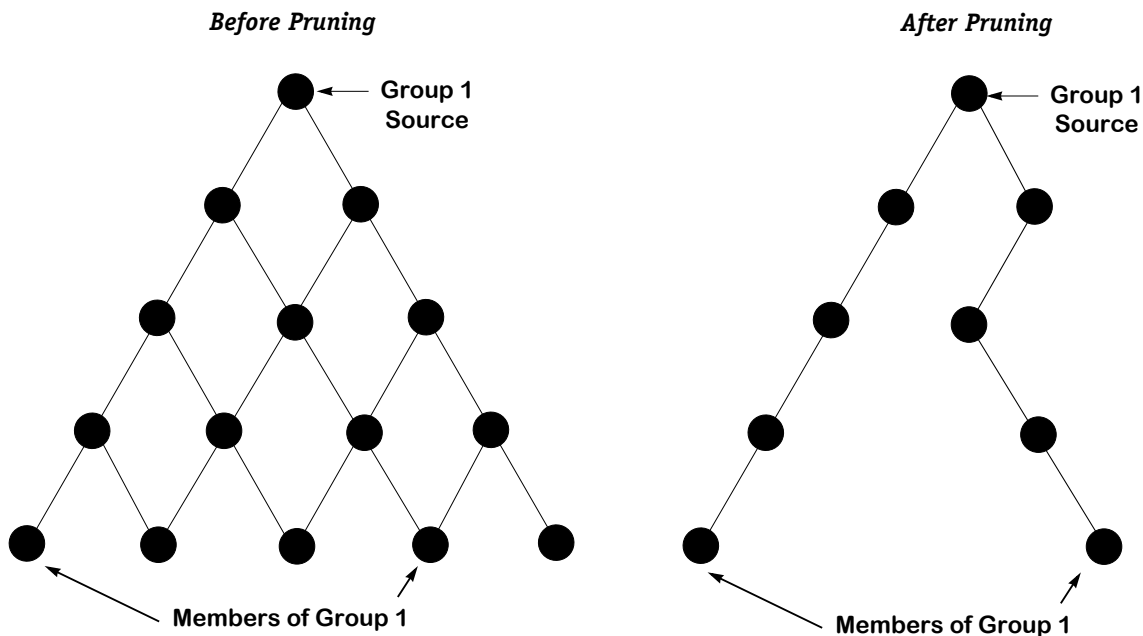


How End Stations Join Multicast Groups

The standard multicast routing application used in Alcatel switches is MrouteD (Multicast Routing Daemon). MrouteD is an implementation of the Distance Vector Multicast Routing Protocol (DVMRP), Version 3, for forwarding multicast datagrams through an internetwork. Unlike the Routing Information Protocol (RIP), which calculates the next hop to a destination address, DVMRP calculates the previous hop back to a source address. Some version of DVMRP is supported in most MBone routers. (See RFC 1075 for more information about DVMRP Version 3.)

Using the Reverse Path Multicasting algorithm, DVMRP constructs a routing table for each multicast group. The routing table, using metrics to determine path cost, determines the shortest path from each multicast host in the group back to the source. When a multicast router receives a multicast packet, it checks its unicast routing tables to make sure the packet arrived on the parent (or expected) interface. If so, it forwards packets out all interfaces that have group membership (except the parent interface) or are connected to other routers. If the packet did not arrive on the parent interface, the packet is discarded.

As traffic is forwarded, routers that do not want traffic send prune messages to signal the upstream router to stop sending traffic. If IGMP membership requests are later received by the router, the router sends graft messages and waits for acknowledgments (graft acks). Eventually the prune times out, and the multicast routing table is reconstructed. The routing table can be described as a routing tree. Branches of the multicast tree that were removed may be grafted into the tree if those stations signal via IGMP that they want to receive the multicast traffic.



Construction of DVMRP Routing Tree

Hardware/Software Support

Multicast routing is supported on OmniSwitches and OmniSwitch/Routers (Omni S/Rs) with at least 16 MB of memory. It is also supported on the following OmniStack models: 2032, 3032, 4024, and 5024.

Multicast traffic is broadcast on all ports of member networks, regardless of whether or not those ports request the multicast traffic. Multicast traffic may be controlled at the switch port level with the IP Multicast Switching feature (IPMS). Note that you should not have any early generation Ethernet modules installed in the switch when IPMS is running. See Chapter 1, “IP Multicast Switching,” for more information about IPMS. For more information about managing Ethernet modules, see your switch user manual.

◆ Note ◆

IP multicast routing is not supported on ATM Classical IP (CIP) connections or WAN interfaces with multiple virtual circuits. Multicast routing also does not support AutoTracker VLANs.

Installing MrouteD

MrouteD is a dynamically loaded module and is contained in the **mrd.img** image file. To upload the file, use standard FTP or ZMODEM procedures. Refer to your switch manual if you need help uploading the software. Typically, IP multicast routing is used in conjunction with IPMS. For information about installing IPMS, see Chapter 1, “IP Multicast Switching.”

Multicast routing is disabled by default and must be enabled through the **mrcfg** command. See *Enabling or Disabling Multicast Routing* on page 2-6.

There is no required configuration for multicast routing. Optional configuration includes disabling pruning and configuring multicast interface parameters. The User Interface (UI) commands are described in the next section.

Disabling MrouteD

You can disable multicast routing through the **mrcfg** command, which is explained in *Enabling or Disabling Multicast Routing* on page 2-6. If you want to disable MrouteD permanently, delete the **mrd.img** file from the flash file system on the switch.

User Interface Commands

When MrouteD is loaded into the switch, a submenu named **IPMR** (for IP multicast routing) is added to the Networking menu of the User Interface (UI).

Commands in the UI are executed by typing the command and pressing **<Enter>**.

◆ **Note** ◆

For general information about the UI, see your switch manual.

IP Multicast Routing Menu

To display the IPMR submenu, enter the following command:

ipmr

If the UI is configured for terse mode, you must then enter a **?** to display the submenu. In verbose mode, the UI automatically displays the submenu.

Command	DVMRP IP Multicast Routing Menu
mrcfg	Configure Multicast Routing as Enabled or Disabled
mrclife	Configure the Time a Cached Multicast Route Stays in Memory
mrftab	Display the Forwarding Cache Table
mrntf	Configure a Multicast Interface
mrprune	Turn Switch into Non-Pruning Router
mrptab	Display Prune Requests Received
mrrtab	Show the Multicast Route Table
mrtables	Show the Multicast Interface Tables
mrtunnel	Configure a Multicast Tunnel
mrupdate	Signal Multicast Router to Use New Configuration

Main

File

Summary

VLAN

Networking

Interface

Security

System

Services

Help

All of the multicast routing commands are described in the remaining sections of this chapter except for the **mrtunnel** command, which is not supported in the current release.

There is no required configuration for multicast routing, and typically default values are adequate. Any configuration change you make requires the **mrupdate** command to activate the change.

Enabling or Disabling Multicast Routing

Multicast routing is disabled by default. Any multicast routing command you enter (except **mrcfg**) is not available when multicast routing is disabled.

◆ Note ◆

Before multicast routing can be enabled using the **mrcfg** command, the **mrd.img** file must be loaded on the switch. See *Installing MrouteD* on page 2-4.

To enable/disable multicast routing:

1. At the system prompt, enter the following command:

```
mrcfg
```

The screen displays similar to the following:

```
Current MrouteD Configuration Status:
```

```
MrouteD status (e=enable, d=disable) : Disabled
```

```
(save/quit/cancel)
```

```
:
```

2. At the prompt, enter **e** for enable or **d** for disable.
3. Enter **save** to save the change, or enter **q** to quit the screen without saving the change or **c** to cancel the change.

Updating the MrouteD Configuration

The MrouteD configuration must be activated to use any configuration changes you make using the **mrclife**, **mrntf**, or **mrprune** command. You can make all of your changes at one time, then enter the **mrupdate** command once to activate all of the changes. Rebooting the switch will also activate any changes you have made.

Typically, after you save a configuration change, the system displays a message reminding you to update the configuration. The message is similar to the following.

Use command 'mrupdate' to activate changes.

To update configuration parameters, enter the following command at the system prompt:

```
mrupdate
```

No message is displayed; the system prompt redisplay.

Viewing Current Cache Life

The “cache lifetime” parameter sets the amount of time, in seconds, that routes are retained in the route cache table.

To display the current value of the cache life parameter, enter the following command:

mrclife

A screen similar to the following displays:

Cache Lifetime is currently set to 300 seconds
Enter a new value or hit return to exit : () :

To change the cache life parameter, enter a new value. For example:

500

A message similar to the following would then display:

Cache Lifetime is now set to 500 seconds

The new value is compared to the allowed range of values (which is 300 to 86,400). If the value entered is valid, the switch configuration file will be updated with the new value. To activate the change, enter the **mrupdate** command as described in the previous section.

Viewing the Multicast Forwarding Table

Use the **mrftab** command to display the source and destination combinations that are currently being forwarded. You can choose to display all the values in the table by entering the **mrftab** command alone. Or, you can use a wildcard search to limit the display to only those addresses that meet a certain criteria. To do so, simply enter a destination IP address after the command, substituting an asterisk for any one of the IP address parameters.

For example, entering

```
mrftab *.*.123.*
```

will display only those addresses with **123** in the third byte. Note that you can specify more than one parameter value. For example:

```
mrftab *.206.123.30
```

If the number of entries in the table exceeds the allocated buffer size, which is determined by the router, the table is truncated and the following message displays:

Too many entries for buffer size. Use wildcards to reduce entries.

To display all the current values in the forwarding cache table, enter the following command:

```
mrftab
```

A screen similar to the following displays:

Multicast Forwarding Table with 1 Entry

Org. IP Addr. =====	Dest. IP Addr. =====	Source =====	Output Groups =====
198.206.10.2	225.0.0.37	2	3
192.203.5.3	225.0.0.6	Not Fwdg	

Fields in the table have the following definitions:

Org. IP Addr. The address from which the multicast datagram originated.

Dest. IP Addr. The IP multicast address to which the datagram is being sent.

Source. Indicates on which VLAN group traffic for the given source is expected to arrive. If datagrams from the specified origin IP address do not arrive on the parent VLAN group, they will be discarded. The number in this field (**2** in the example) corresponds to the number in the **Grp** field in the group table (see *Viewing the Group Table* on page 2-20). If **Not Fwdg** is displayed, the router has not established the source group because it is exchanging routing messages with other routers, or the source group has been disabled by **mrintf** command.

Output Groups. Indicates on which VLAN group(s) the datagram will be forwarded (which may be workstations as well as other routers). The number in this field (**3** in the example) corresponds to the number in the **Grp** field in the group table (see *Viewing the Group Table* on page 2-20).

Displaying the Interface Configuration Table

Use the **mrintf** command to display, add entries to, or update the interface configuration table. When MrouteD is started, all interfaces are discovered during the initialization process and assigned default values. You will only need to make additions to the interface table if you want to override one or more of the default values.

1. To display the interface configuration table, enter the following command:

mrintf

A message similar to the following displays:

Do you wish to see the list of current physical interface entries? (Yes) :

2. Press **<Enter>** to display the contents of the table.

A screen similar to the following displays:

```

Multicast Interface table
Index  I/F Address  Enable  Silent  Ver  Metric  Thresh  Rate limit  Bound  Alt_nets
=====
No entries
To modify an interface, type the index of the corresponding entry
To add type a, delete type d, exit type e (e) :

```

Fields in the table have the following definitions:

Index

An index reference used to select an entry for updating.

I/F Address

The IP address of the interface.

Enable

Indicates whether or not the interface is enabled or disabled.

Silent

Indicates whether or not (**Y** or **N**) the interface will send or accept neighbor probes. If the interface is configured not to send or accept probes, it will receive only local multicast traffic.

Ver

The IGMP version (**1** or **2**).

Metric

The weight or cost of this connection as compared to other connections on the switch. If a connection has a slower media speed, or a greater number of hops, the configured metric should be larger. Acceptable values are 1 to 31.

Threshold

The number of hops or Time to Live (TTL) that a packet needs to contain before it will be forwarded on the interface. For example, if the TTL of a datagram is 30 but the threshold is configured to be 64, the packet will not be forwarded. Acceptable values are 1 to 255.

Rate Limit

The maximum data rate, in kilobits per second, that this interface can use for multicast traffic. Acceptable values are 0 to 100,000. (A value of 0 means that the rate is unlimited.)
When IPMS is running on the switch, this value has no effect.

Bound

Indicates whether or not boundary tables are associated with the interface. Each boundary table includes the scoped address and a mask length. The scoped address must have 239 as the first byte of the IP address. When a user on the specified interface joins the multicast group defined by the scoped address plus the mask length, all multicast traffic will stop being forwarded on the interface. This provides a mechanism for the end user to control multicast traffic from the network.

Alt_nets

Indicates whether or not alternate network tables are associated with the interface. Alternate network tables are used to declare routers that may be located on the same network as the router defined by the primary address on the virtual interface table. If a router on the primary network is supposed to receive a control message, another copy of the message is generated for each of the alternate networks associated with the primary network. The alternate network tables contain the network address and the subnet mask length.

Adding an Entry to the Interface Configuration Table

This section shows an example of adding an entry to the interface configuration table. This procedure describes a sample configuration. Values given are examples *only* and are not required or recommended.

1. Use the **mrntf** command to display the interface configuration table. See *Displaying the Interface Configuration Table* on page 2-9 for information about the parameters in this table.

2. Add an entry to the multicast interface table by entering **a** at the following prompt:

To add type a, delete type d, exit type e (e) :

The screen displays the following prompt:

IP interface: (x.x.x.x) (0.0.0.0)?

3. Enter an IP interface in dotted-decimal format. For example:

198.206.10.1

4. The screen displays a series of prompts. For this example, press **<Enter>** after each prompt to accept the default values:

Enable/Disable Multicast (e=enabled, d=disabled) (ENABLED)?
Select Silent Mode (s=silent, n=normal) (NORMAL)?
Metric: (1-31) (1)?
Threshold: (1-255) (1)?
Rate Limit (kilobits/sec): (0-100000) (500)?
IGMP Version: (1=V1, 2=V2) (2)?

The screen displays the following prompt:

Boundary List
Do you wish to see the list of current physical boundary entries? (Yes) :

5. Press <Enter> to display the contents of the table. A message similar to the following displays:

No entries

To modify a boundary, type the index of the corresponding entry

To add type a, delete type d, exit type e (e) :

6. Enter **a** to indicate that you want to add an entry. The screen displays the following prompt:

Scoped Address: (239.x.x.x)

A scoped address must be a valid IP address with 239 in the first byte. For more information, see the definition of the Bound parameter in *Displaying the Interface Configuration Table* on page 2-9.

7. Enter an IP address. For example:

239.1.1.0

The screen displays the following prompt:

Mask Length: (1-31)

The mask length is added to the Scoped Address to create a boundary. When a user on a specified network joins the multicast group defined by this boundary, no multicast traffic can be forwarded on the network. See the definition of the Bound parameter in *Displaying the Interface Configuration Table* on page 2-9.

8. Enter **24** for the mask length. The screen displays the following prompt:

To modify a boundary, type the index of the corresponding entry

To add type a, delete type d, exit type e (e) :

9. Enter **e** to indicate that you want to exit. The screen displays the following prompt:

Alt Net List

Do you wish to see the list of current alternate net entries? (Yes) :

10. Press <Enter> to display the contents of the table. A screen similar to the following displays:

No entries

To modify an altnet, type the index of the corresponding entry

To add type a, delete type d, exit type e (e) :

11. Enter **a** to indicate that you want to add an entry. The screen displays the following prompt:

Alternate Net Address: (x.x.x.x)

An alternate network address indicates that a device on the specified network is using a different IP address. See the definition of the Alt_nets parameter in *Displaying the Interface Configuration Table* on page 2-9.

12. Enter an IP address in dotted-decimal format. For example:

198.206.18.0

The screen displays the following prompt:

Mask Length: (1-31)

13. Enter **24** for the mask length. The screen displays the following prompt:

**To modify an altnet, type the index of the corresponding entry
To add type a, delete type d, exit type e (e) :**

14. Enter **e** to indicate that you want to exit. The screen displays the following prompt:

Do you wish to save your changes? (Yes) :

15. Press **<Enter>** to save the changes. The screen displays the following message:

**Saved
message?????**

16. Use the **mrupdate** command to activate the configuration (see “Updating MrouteD Configuration” on page 2-6).

Modifying an Entry in the Interface Configuration Table

The following procedure describes how to modify an existing entry.

1. Use the **mrintf** command to display the interface configuration table. See *Displaying the Interface Configuration Table* on page 2-9 for information about the parameters in this table.

A sample table is shown here.

```

Multicast Interface table
Index I/F Address Enable Silent Ver Metric Thresh Rate limit Bound Alt_nets
=====
1      198.206.10.1 E      Y    2    1    1      500 KBPS   Yes   Yes
To modify an interface, type the index of the corresponding entry
To add type a, delete type d, exit type e (e) :

```

2. Enter the number of the index entry that you want to modify (in this example, **1**). The screen displays the following table and prompt:

```

Current values of this interface are:
1) IP interface: (x.x.x.x) 198.206.10.1
2) Enable/Disable Multicast (e=enabled, d=disabled) ENABLED
3) Silent mode (s=silent, n=normal) NORMAL
4) Metric: (1-31) 1
5) Threshold: (1-255) 1
6) Rate Limit (kilobits/sec): (0-100000) 500
7) IGMP Version: (1=V1, 2=V2) 2
8) Boundary List
9) Alt Net List
(quit/cancel or enter index = value to update)
:

```

The fields are described in *Displaying the Interface Configuration Table* on page 2-9.

3. Make changes to the fields 1 through 7 by entering the number of the relevant parameter, an equal sign, and the desired value. For example, enter **6=800** to indicate that you want to change the Rate Limit field to 800. The screen redisplay as follows:

```

Current values of this interface are:
1) IP interface: (x.x.x.x) 198.206.10.1
2) Enable/Disable Multicast (e=enabled, d=disabled) ENABLED
3) Silent mode (s=silent, n=normal) NORMAL
4) Metric: (1-31) 1
5) Threshold: (1-255) 1
6) Rate Limit (kilobits/sec): (0-100000) 800
7) IGMP Version: (1=V1, 2=V2) 2
8) Boundary List
9) Alt Net List
(quit/cancel or enter index = value to update)

```

The boundary list and the alternate network list are available from fields 8 and 9 (see *Modifying the Boundary List* on page 2-14 and *Modifying the Alternate Network List* on page 2-15).

4. Enter **q** to exit this screen. The screen displays the following prompt:

Do you wish to save your changes? (Yes) :

5. Press **<Enter>** to save the changes. The screen displays the following message:

Saved

6. Use the **mrupdate** command to activate the configuration (see “Updating MrouteD Configuration” on page 2-6).

Modifying the Boundary List

This section describes how to modify the boundary list. A boundary is a multicast group defined by an IP address called a scoped address (in the range 239.0.0.0 to 239.255.255.255), and a subnet mask. The boundary prevents any multicast traffic from reaching any network in which there is at least one end station that has requested membership in the multicast group.

1. Display the current values of the interface configuration table as described in *Modifying an Entry in the Interface Configuration Table* on page 2-13.

2. Enter **8** at the **mrntf** command prompt. The screen displays the following message:

Do you wish to see the list of current physical boundary entries? (Yes) :

3. Press **<Enter>** to display the contents of the table. A screen similar to the following displays:

```
Boundaries:
Index  Scoped Address  Mask Length
=====
  1      239.1.1.0      24
```

**To modify a boundary, type the index of the corresponding entry
To add type a, delete type d, exit type e (e) :**

Entries may be added or deleted from the table. This example shows how to modify an existing entry that was configured in *Adding an Entry to the Interface Configuration Table* on page 2-10. For information about boundary parameters, see the same section.

4. Enter the number of the boundary that you want to change (in this example, **1**). A screen similar to the following displays:

Current values of this boundary entry are:

1) Scoped Address: (239.x.x.x)239.1.1.0

2) Mask Length: (1-31) 24

(quit/cancel or index = value to update)

:

5. Enter the value of the parameter that you want to change, followed by an equal sign and the desired value. For example:

1=239.2.2.0

The boundary entry redisplay with the new value. Make any change to the Mask Length in the same way.

6. Enter **q** to exit the boundary parameters. The boundary table displays.

- Enter **e** to indicate that you want to exit. The screen displays the current interface values and prompt:

```

Current values of this interface are:
1) IP interface: (x.x.x.x) 198.206.10.1
2) Enable/Disable Multicast (e=enabled, d=disabled) ENABLED
3) Silent mode (s=silent, n=normal) NORMAL
4) Metric: (1-31) 1
5) Threshold: (1-255) 1
6) Rate Limit (kilobits/sec): (0-100000) 800
7) IGMP Version: (1=V1, 2=V2) 2
8) Boundary List
9) Alt Net List
(quit/cancel or enter index = value to update)
:

```

- Enter **q** to exit this screen. The screen displays the following prompt:

```
Do you wish to save your changes? (Yes) :
```

- Press **<Enter>** to save the changes. The screen displays the following message:

```
Saved
```

- Use the **mrupdate** command to activate the configuration (see “Updating MrouteD Configuration” on page 2-6).

Modifying the Alternate Network List

This section describes how to modify the alternate network list. When a device on the network uses an IP address that is different from the network address, an alternate network address must be defined so that device may receive multicast traffic.

- Display the current values of the interface configuration table as described in *Modifying an Entry in the Interface Configuration Table* on page 2-13.
- Enter **9** at the **mrntf** command prompt to indicate that you want to examine the alternate network list. The screen displays the following prompt:

```
Do you wish to see the list of current alternate net entries? (Yes) :
```

- Press **<Enter>** to display the contents of the table. The screen displays the following table and prompt:

```

Alternate Networks:
Index  Scoped Address  Mask Length
=====
1      198.206.18.0    24
To modify an altnet, type the index of the corresponding entry
To add type a, delete type d, exit type e (e) :

```

Entries may be added or deleted from the table. This example shows how to modify an existing entry that was configured in *Adding an Entry to the Interface Configuration Table* on page 2-10. For information about parameters in the alternate network table, see the same section.

4. Enter the number of the alternate network that you want to change (in this example, **1**). A screen similar to the following displays:

```
Current values of this boundary entry are:  
1) Scoped Address: (239.x.x.x)198.206.18.0  
2) Mask Length: (1-31) 24  
(quit/cancel or index = value to update)  
:
```

5. Enter the value of the parameter that you want to change, followed by an equal sign and the desired value. For example:

```
1=198.206.20.0
```

The alternate network entry redisplay with the new value. Make any change to the mask length in the same way.

6. Enter **q** to exit the alternate network parameters. The alternate network table displays.
7. Enter an **e** to indicate that you want to exit. The screen displays the current interface values and prompt:

```
Current values of this interface are:  
1) IP interface: (x.x.x.x) 198.206.10.1  
2) Enable/Disable Multicast (e=enabled, d=disabled) ENABLED  
3) Silent mode (s=silent, n=normal) NORMAL  
4) Metric: (1-31) 1  
5) Threshold: (1-255) 1  
6) Rate Limit (kilobits/sec): (0-100000) 800  
7) IGMP Version: (1=V1, 2=V2) 2  
8) Boundary List  
9) Alt Net List  
(quit/cancel or enter index = value to update)  
:
```

8. Enter **q** to indicate that you want to quit. The screen displays the following prompt:

```
Do you wish to save your changes? (Yes) :
```

9. Press **<Enter>** to save the changes. The screen displays the following message:

```
Saved
```

10. Use the **mrupdate** command to activate the configuration (see “Updating MrouteD Configuration” on page 2-6).

Disabling or Enabling Pruning

Use the **mrprune** command to disable or enable pruning in the multicast router. A pruning router will send a prune message to the upstream router if there are no members for the received datagram. For more details on pruning, see *Introduction* on page 2-1. Pruning is enabled by default. If you set pruning to disabled, the upstream router will continue to forward traffic, and the receiving router will simply discard all datagrams for which there is no demand.

To disable pruning, follow the steps here:

1. Enter the following command:

```
mrprune
```

A screen similar to the following displays:

```
Pruning is currently set to Enabled
To enable type e, to disable type d, hit return to exit : () :
```

2. Enter **d** at the prompt. A message displays similar to the following:

```
Pruning is now set to Disabled
```

3. Use the **mrupdate** command to activate the configuration. See “Updating MrouteD Configuration” on page 2-6.

To enable pruning again, enter the **mrprune** command, enter **e** for enable, and enter the **mrupdate** command. A configuration message displays.

Viewing Prune Requests

Prune requests are sent from a neighboring router to indicate that it has not received any membership reports for a particular multicast group or it does not have any downstream neighbors with membership in a particular multicast group.

Use the **mrptab** command to view a list of all prune requests received by this switch. To display only those prune requests for a particular multicast group, include the multicast group address with the command. The address may include wildcards (asterisks). For example, if you want to display the prune requests for multicast groups with 225 in the first byte, you would enter:

```
mrptab 225.*.*
```

A screen similar to the following displays:

MrouteD Current Prune Requests

Current number of Prune Requests Displayed: 2

Multicast Group	Source IP Network	Pruning Router	Rtr Group	Timer
=====	=====	=====	=====	=====
225.1.1.1	13.1.100/22	198.28.5.70	1	290
225.2.2.2	13.1.100/22	198.28.5.70	1	298

The fields are defined as follows:

Multicast Group. The IP address of the multicast group, which is associated with a multicast stream. This address is sometimes referred to as the destination address.

Source IP Network. The IP address and the subnet mask of the source network from which the multicast stream originated.

Pruning Router. The IP address of the router that is sending the prune messages. Pruning messages instruct MrouteD to stop sending traffic for the multicast group to the pruning router.

Rtr Group. The VLAN group associated with the router that is sending the prune messages.

Timer. The amount of time (in seconds) until the prune times out. If the timer expires and no additional prune messages have been received from the pruning router, multicast traffic will start flowing to the router.

Viewing the Multicast Routing Table

Use the **mrrtab** command to display a list of all the routes currently in the route table. To limit the display to a particular source network (or origin-subnet), specify the network IP address, for example:

```
mrrtab 192.168.10.0
```

You can use a wildcard search to limit the display to particular origin-subnet addresses. For example, entering

```
mrrtab *.*.123.*
```

displays only those addresses that have **123** in the third byte.

To display a list of all the routes currently in the route table, enter the following command:

```
mrrtab
```

A screen similar to the following displays:

Multicast Routing Table with 7 Entries

Origin-Subnet	From-Gateway	Metric	Tmr	In-Grp	Out-Groups
=====	=====	=====	=====	=====	=====
198.206.184/24		1	50	2	1* 3* 4* 5*
198.206.58/24		1	50	5	1* 2* 4 5
198.206.11/24		1	50	4	1* 2* 3 5*
198.206.10/24		1	50	3	1* 2* 4 5*
198.206.7/24	198.206.11.2	2	30	4	1* 2* 3 5*
198.168.10/24	198.206.10.2	2	10	3	1* 2* 4 5*
192.168.10/24		1	50	1	2* 3 4 5*

The screen header indicates the total number of entries in the table. Fields in the table have the following definitions:

Origin-Subnet. The IP address of the network where the datagram originates. Also referred to as the source network.

From-Gateway. The network IP address of the network on which the datagram enters the router.

Metric. The weight or cost of this connection as compared to other connections on the switch. If a connection has a slower media speed, or a greater number of hops, the configured metric should be larger. The range of values is 1 to 31 and is configured on the multicast interface table (see *Displaying the Interface Configuration Table* on page 2-9).

Tmr. The timer value for this route. The timer is incremented every 5 seconds. When the timer reaches the route timeout value (140 seconds), the route is deleted. If a route report is received for this route, the timer is reset to zero.

In-Grp. The VLAN group number of the sending interface.

Out-Groups. The VLAN groups that will receive outgoing traffic for the specified origin-subnet and source gateway. Forwarded traffic will be sent on all listed VLAN groups except leaf destinations. An asterisk (*) next to an interface number indicates that the group is a leaf. (A leaf group is one that has no other routers connected to it.)

Viewing the Group Table

Use the **mrtables** command to display the VLAN group table, which contains information about VLAN groups on the system. The table contains one entry for each group defined in the switch.

To limit the display to a particular VLAN group, enter the **mrtables** command with the desired group number. For example:

```
mrtables 3
```

To display all entries in the group table, enter the command without a group number:

```
mrtables
```

A screen similar to the following displays:

```
Mrouted version: 3.8. Switch Version: 4.1.2.0 Release Build Jan 28 2000 at 08:45:12
```

```
Groups with neighbors = 1
[This host is a leaf]
```

```
Group Table
Grp      Local Address  Interface Information  M  Thr  Rate  Flags
=====  =====
1        172.28.5.3     subnet: 172.28.5/24    1   1    0    querier leaf
                        peers: 172.28.5.86 (11.0) (0xa) 30 seconds
                        pkts in : 0
                        pkts out: 0

2        193.166.10.3  subnet: 193.166.10/24  1   1    0    querier leaf
                        pkts in : 0
                        pkts out: 0

3        193.180.10.13 subnet: 193.180.10/24  1   1    0    querier leaf
                        pkts in : 0
                        pkts out: 0

4        13.1.103.3    subnet: 13.2.16/22     1   1    0    querier leaf
                        groups: 225.2.2.2
                        pkts in : 1029
                        pkts out: 0
```

The screen header shows the MrouteD version and the number of groups with neighbors. In this example, only one group (group 1) has a neighbor (peer). The fields in the table have the following definitions:

Grp. The VLAN group number associated with the local address. The multicast forwarding table and the multicast routing table refer to this field (see *Viewing the Multicast Forwarding Table* on page 2-8 and *Viewing the Multicast Routing Table* on page 2-19).

Local Address. The IP address assigned to this virtual interface.

Interface Information. The first line in this column will contain one of two possible entries. If the connection is a local connection, the first line of the field will display the IP address and subnet mask (for example, **subnet: 192.168.10/24**). Immediately under the first line any one of the subentries listed below may display, depending on the network topology and configuration:

alternate subnets

Displays the IP address and subnet mask if there are alternate subnets configured on the multicast interface table (see *Displaying the Interface Configuration Table* on page 2-9 for information about configuring alternate subnets).

peers

Displays the IP address of the peer if the virtual interface has established a peer or neighbor relationship with an adjacent router. The MrouteD version the peer is running displays in parentheses, followed by a hexadecimal representation of the peer's flags. Also displays the number of seconds since the last neighbor probe was received.

groups

Displays the destination address (multicast group address) if users have joined groups on the virtual interface. Also displays the number of seconds before the group times out waiting to receive a neighbor probe. The timer is reset to 260 seconds after it times out.

boundaries

Displays the scoped address and subnet mask if any boundary tables have been configured for this virtual interface (for information about configuring boundary tables, see *Displaying the Interface Configuration Table* on page 2-9).

pkts in

The number of packets received on this interface.

pkts out

The number of packets sent out on this interface.

M. Indicates the metric, which is the weight or cost of this connection as compared to other connections on the switch. If a connection has a slower media speed, or a greater number of hops, the configured metric should be larger. The range of values is 1 to 31.

Thr. Indicates the number of hops or Time to Live (TTL) that a packet needs to contain before it will be forwarded. For example, if the TTL of a datagram is 30 but the threshold is configured to be 64, the packet will not be forwarded. The range of values is 1 to 255.

Rate. Indicates the maximum data rate, in kilobits per second, that can be sent on this interface. Acceptable values are 0 to 100,000 (a value of 0 means that the rate is unlimited). (*Note this value is not used if IPMS is running on the switch.*)

Flags. Indicates where the flags in the interface table are analyzed and displayed. The values are as follows:

querier	The interface is the network querier (peer with the lowest IP address).
one-way	A one-way interface (sending probes but not receiving probes).
down	Kernel status of the interface.
disabled	Administratively disabled.
leaf	Whether or not the peer is a leaf router.
IGMPv1	This interface is to act as an IGMPv1 interface.
silent	The interface does not send or respond to neighbor probes.

