

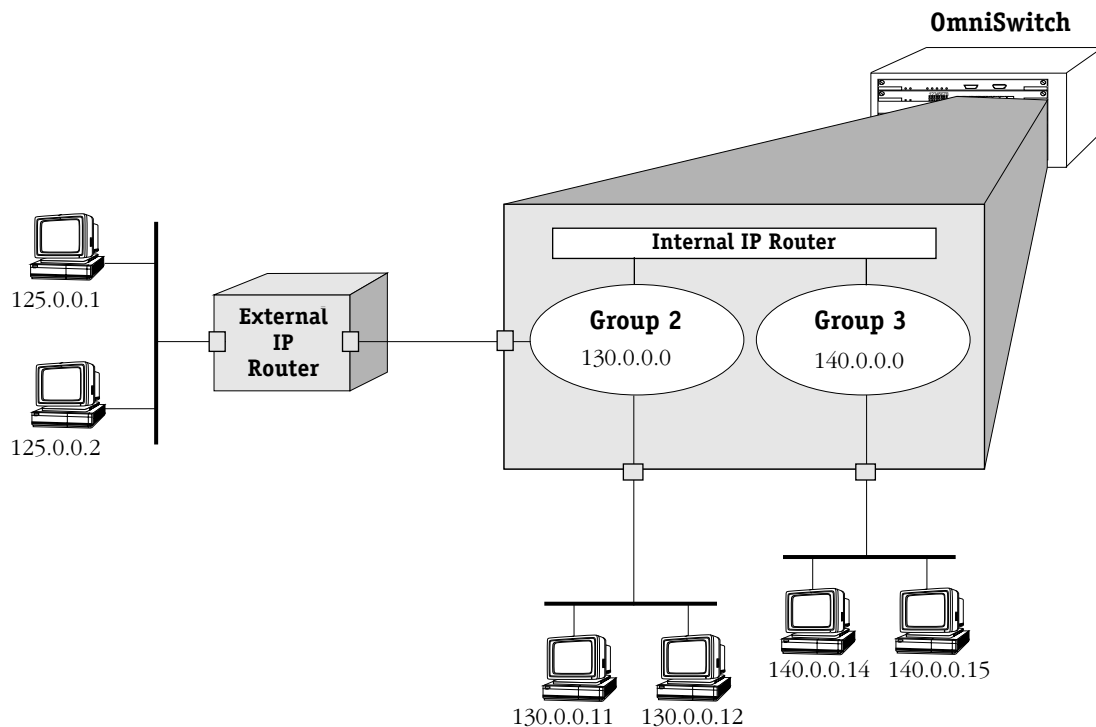
31 IP Routing

Introduction

This chapter gives an overview of IP routing and includes information about configuring static routes and viewing/configuring TCP/IP protocols such as Telnet and the Routing Information Protocol (RIP). IP routing requires at least one virtual router port to be configured on the switch. For information about configuring virtual router ports, see Chapter 25, “Managing Groups and Ports.”

When IP routing is enabled on the switch, the switch will be able to exchange routing information with external IP routers in the network, and stations connected to groups and VLANs with virtual router ports will be able to communicate. Groups or VLANs that do not have router ports with routing enabled are essentially firewalled from each other.

In the example shown here, stations connected to each group will be able to communicate if a virtual router port is created for each group and each router port on the switch has IP routing enabled. Stations in group 2 and group 3 will also be able to communicate with stations attached to the external IP router if a default route to that router is configured on the switch or the switch learns about the external router through RIP or some other routing protocol.



IP Routing Overview

In switching, traffic may be transmitted from one media type to another within the same broadcast domain (or group/VLAN). Switching happens at layer 2, the physical layer; routing happens at layer 3, the network layer. In routing, traffic may be transmitted across groups/VLANs, and broadcast or multicast traffic is prevented from being transmitted across those domains (unless some other mechanism is set up on the switch, such as UDP forwarding or IP multicast routing).

In IP routing, the switch builds routing tables to keep track of optimal destinations for traffic it receives that is destined for remote networks. The switch also sends and receives routing messages, or advertisements, to/from other routers in the network. When the switch receives a packet to be routed, it strips off the MAC header and examines the IP header of the packet. It looks up the source/destination address in the routing table, and then adds the appropriate MAC address to the packet.

Calculating routing tables and stripping/adding MAC headers to packets is performed by switch software unless a Hardware Routing Engine (HRE) or HRE-X is installed. The HRE or HRE-X significantly improves routing performance. See Chapter 1, “Omni Switch/Router Chassis and Power Supplies,” and Chapter 6, “The MPM,” for information about the HRE-X and HRE respectively. For the Omni S/R, IP routing speed is boosted by a fastpath mechanism included in the base code (see *Viewing IP Statistics and Errors* on page 31-12).

IP is associated with several layer 3 and layer 4 protocols. Some of these protocols are built into the base code loaded into the switch. Others are included as part of Advanced Routing software. Some protocols are specifically used for routing; others are used by any host or end station that has an IP address. A brief overview of supported IP protocols is included here.

Routing Protocols

When IP routing is enabled, the switch uses routing protocols to build routing tables that keep track of stations in the network and to decide the best path for forwarding data. These routing protocols include:

- **Routing Information Protocol (RIP)**—An interior gateway protocol that defines how routers exchange information in an autonomous system. RIP makes routing decisions using a “least-cost path” method. RIP services are performed by a program operating in the switch called RouteD. RIP and RIP II services are also available from a program called GateD, which is part of Alcatel’s optional Advanced Routing software. RIP, whether performed by RouteD or GateD, allows the switch to learn routing information from other, neighboring RIP routers.
- **Open Shortest Path First (OSPF)**—An interior gateway protocol that provides a routing function similar to RIP but which uses different techniques to determine the best route for a datagram. OSPF services are provided by GateD, part of Alcatel’s optional Advanced Routing software.
- **Border Gateway Protocol (BGP)**—An exterior gateway protocol that provides for routing between autonomous systems. BGP is not part of the base code but is included in the Advanced Routing software.

Transport Protocols

IP is both connectionless (it routes each datagram separately) and unreliable (it does not guarantee delivery of datagrams). This means that a datagram may be damaged in transit, or thrown away by a busy router, or simply never make it to its destination. The resolution of these transit problems is to use a layer 4 transport protocol:

- Transmission Control Protocol (TCP)—A major data transport mechanism that provides reliable, connection-oriented, full-duplex data streams. While the role of TCP is to add reliability to IP, TCP relies upon IP to do the actual delivering of datagrams.
- User Datagram Protocol (UDP)—A secondary transport-layer protocol that uses IP for delivery. However, UDP is not connection-oriented so it does not provide reliable end-to-end delivery of datagrams. But some applications can safely use UDP to send datagrams that don't require the extra overhead added by TCP.

Application-Layer Protocols

- Bootstrap Protocol (BOOTP)/Dynamic Host Configuration Protocol (DHCP)—May be used by an end station to obtain an IP address. The switch provides a UDP relay that allows BOOTP requests/replies to cross different networks. See Chapter 32, “UDP Forwarding.”
- Simple Network Management Protocol (SNMP)—Used to manage nodes on a network. SNMP is discussed in Chapter 17, “Configuring SNMP.”
- Telnet—Used for remote connection to a device. The **telnet** command is described in this chapter.
- File Transfer Protocol (FTP)—Enables transferring files between hosts.

Additional IP Protocols

- Internet Control Message Protocol (ICMP)—Specifies the generation of error messages, test packets, and informational messages related to IP. ICMP supports the **ping** command used to determine if hosts are online.
- Address Resolution Protocol (ARP)—Used to find the IP address that corresponds to a given physical (MAC) address.
- Internet Group Management Protocol (IGMP)—Tracks multicast group membership. See the Multicast Services section of the *Advanced Routing User Manual*.
- Resource ReSerVation Protocol (RSVP)—Signals Quality of Service (QoS) requests in an IP network. For more information, see the *Switched Network Services User Manual*.

Setting Up IP Routing on the Switch

IP routing is enabled on a per-port basis by creating a virtual IP router port for a group/VLAN. The switch does not do any routing unless the virtual router port has IP routing enabled (routing is enabled by default). The steps for setting up IP routing on the switch are given here:

Step 1. Configuring a Virtual Router Port

A virtual router port may be created when you set up or modify a group/VLAN through the **crnp** command or **modvl** command described in Chapter 25, “Managing Groups and Virtual Ports.” To create a virtual router port, you enable IP routing and specify an IP address for the router port.

When routing is enabled on the port, the switch creates routing tables and address translation tables so it knows how to forward traffic. The switch keeps track of router ports and any other routers in the network. The switch uses the Address Resolution Protocol (ARP) to match IP addresses with MAC addresses. It uses routing protocols, such as the Routing Information Protocol (RIP), to determine the best path for forwarding traffic. (Other routing protocols are available in the Advanced Routing software package.) It also periodically sends/receives routing messages to/from other routers to keep its routing tables updated.

◆ Important Note ◆

When Spanning Tree and IP routing are both enabled, packets will not be forwarded unless the Spanning Tree Status for the port to which packets are to be forwarded has progressed from Listening to Learning to Forwarding. For example, if IP is enabled on VLAN 42 that has ports 1/1-3 attached to it and you want to forward to a host from port 1/2. You can use the **vi 1/2** command to determine if the Spanning Tree Protocol has entered the Forwarding state for that port.

Step 2. Configuring Optional IP Routing Parameters

Optional configuration for IP routing includes the following:

- Static routes. These are routes that are manually added to the routing table and may be used rather than dynamic routes (which are learned through routing protocols like RIP).
- RIP filters. Controls the operation of RIP by minimizing the number of entries that will be added to the routing table.

Static routes and RIP filters are described in this chapter. This chapter also describes how to view various IP statistics as well as the routing table. It includes information about how to ping another IP host in the network to see if you have connectivity, how to telnet to a remote system, or trace an IP route.

Step 3. Configuring Other IP Routing Features

There are several optional features that may be used with IP routing. Some features are included as part of the base code and are described in this user manual. Other features are available as optional switch software and are described in separate user manuals. The features are listed here:

- UDP forwarding—Forwards UDP broadcasts/multicasts across groups/VLANs. See Chapter 32, “UDP Forwarding.”
- HRE-X filtering—Filters IP traffic at the router port if there is an HRE-X installed. See Chapter 34, “HRE-X Filtering.”
- GateD—Provides gateway protocols, including RIP, OSPF, and BGP/CIDR. See the *Advanced Routing User Manual*.
- Virtual Router Redundancy Protocol (VRRP)—Used to back up static IP routes. See the *Advanced Routing User Manual*.
- IP Firewall—Enables the switch to act as a gateway to provide security for all data entering and exiting the switch to and from its attached physical ports, as well as internally between groups and VLANs that are defined in the switch. See the *Switched Network Services User Manual*.
- Multicast services—Includes IP multicast switching (IPMS) and IP multicast routing (MrouteD). See the *Advanced Routing User Manual*.
- IP Control—Manages IP addresses through Lightweight Directory Access Protocol (LDAP), DHCP, and Domain Name Service (DNS). See the *Switched Network Services User Manual*.

The Networking Menu

The Networking menu contains commands that control, and are related to, the routing protocols that are run on the switch.

To switch to, and to display, the **Networking** menu, enter the following commands:

```
networking
?
```

If you have enabled the verbose mode, you do not need to enter the question mark (?).

A screen similar to the following displays:

Command	Networking Menu
snmps	View SNMP statistics
snmpc	Configure SNMP
Names	Configure the DNS resolver
probes	Display all RMON probes
events	Display all logged RMON events
IP	Enter IP networking command sub-menu.
IPX	Enter IPX networking command sub-menu
Gated	Enter Gated menu/control Gated
IPMR	Enter the IPMR routing sub-menu
IPMS	Enter the IPMS networking command sub-menu
VRRP	Enter the VRRP menu
QoS	Enter the QoS menu
Policy	Administer the SNS policy sub-menu
LDAP	Configure the SNS LDAP server sub-menu
pm	Enter port monitor utility command sub-menu
chnghmac	Configure router port's MAC address on selected Group

Main	File	Summary	VLAN	Networking
Interface	Security	System	Services	Help

The commands in this menu are described throughout this manual as follows:

- The **snmps** and **snmpc** commands are described in Chapter 17, “Configuring SNMP.”
- The **Names**, **probes**, **events**, and **chnghmac** commands are described in Chapter 18, “RMON and DNS Resolver.”
- The IP submenu is discussed in this chapter. The IPX submenu is described in Chapter 33, “IPX Routing.”
- The Gated, IPMR, IPMS, and VRRP submenus are available if you have Advanced Routing software loaded. See the *Advanced Routing User Manual* for more information.
- The QoS, Policy, and LDAP submenus are available if you have Switched Network Services software loaded. See the *Switched Network Services User Manual* for more information.
- The **pm** command is described in Chapter 25, “Managing Groups and Ports.”

The IP Submenu

The **ip** command in the Networking menu is used to display the IP submenu. To display the IP submenu, enter the following commands:

```
ip
```

```
?
```

If you have enabled the verbose mode, you don't need to enter the question mark (?).

A screen similar to the following displays:

Command	IP Menu
xlat	View the address translation table
ips	View IP stats & errors
ipr	View IP routes
aisr	Add an IP static route
risr	Remove an IP static route
icmps	View ICMP stats & errors
ping	Ping a system
udps	View UDP stats and errors
udpl	View the UDP listener table
rips	View RIP stats and errors
tcps	View TCP-related statistics
tcpc	View the TCP Connection table
telnet	Remote login to another system using TELNET
tracroute	Trace an IP route
relay	Use 'relayc' or 'relays'
fwconfig	Configure the IP Firewall
ripflush	Flush all routes obtained by RIP
ipfilter	Add/delete an IP RIP filter
ipf	Display IP RIP filters
ipmac	View the IP to MAC Address Association table
ipclass	Turn on/off IP Class Address Checking
Main Interface	File Security
	Summary System
	VLAN Services
	Networking Help

This chapter describes all of the above commands with the exception of **fwconfig**, **relayc**, **relays**, and **ipclass** commands. The **fwconfig** command is described in the *Switched Network Services User Manual*. The relay commands, **relayc** and **relays**, are described in Chapter 32, "UDP Forwarding." The **ipclass** command is described in the *Advanced Routing User Manual*.

The remaining sections of this chapter cover the above commands in the order in which they display in the IP submenu.

Viewing the Address Translation (ARP) Table

The **xlat** command is used to access the ARP (Address Resolution Protocol) Table. This table contains a listing of IP addresses and their corresponding translations to MAC addresses (or slot/port for WAN interfaces). Submenu commands are used to add entries to the table, to delete them, show all the entries currently in the table, to flush “temporary” entries, to display specific entries by either MAC or IP address, and to quit out of the **xlat** submenu.

To begin working with the ARP Table, enter the following command:

xlat

A screen similar to the following displays:

ARP Table Functions

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

The default command is **show** which is used to display all entries in the table. The **quit** command is used to exit out of this submenu and return to the main system prompt.

Displaying All Entries in the ARP Table

At the above prompt, press **<Enter>** to select **Show**, the default command.

A screen similar to the following displays:

Address Translation Table

IP Address	at	Physical Address
90.0.0.1	at	3/1, dlci=32
198.206.184.34	at	00:05:02:c0:7f:11
198.206.184.254	at	00:20:da:6a:98:40

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

The fields on this screen have the following meanings:

IP Address

The IP address, in dotted-decimal format, of a specific host or other device.

Physical Address

The MAC address, in hexadecimal format, of the specific host or other device that corresponds to the IP address in the left-hand column.

Adding Entries to the ARP Table

The **add** subcommand is used to manually add an IP address entry to the ARP Table. To be able to manage your switch over an IP network connection, you will need at least one IP address configured for the switch.

Follow the steps below to add an address to the ARP Table.

1. Enter **add**.

The following prompt displays:

Host name or IP addr to add:

Enter the name of the host or its IP address.

2. The following prompt displays:

Physical address (format aa:bb:cc:dd:ee:ff):

Enter the host's physical address in hexadecimal format.

3. The following prompt displays:

Publish (i.e., proxy for) this entry? (y/n) (n):

Enter **y** to publish (i.e., proxy for) this ARP entry. This feature allows the switch to answer all ARP requests directed at the hosts on a subnetwork. As the "proxy" for these hosts, the switch responds with its own MAC address whenever ARP requests come in for any of the hosts on the subnetwork. Enter **n** if you do not want this ARP entry to act as a proxy.

4. The following prompt displays:

Is this entry permanent (ie. flush will not remove it) (y/n)? (n) :

Enter **y** if this entry is to be permanent (that is, you do not want it to be removed by the **Flush** subcommand). Enter **n** if the entry is to be temporary (that is, you want to allow it to be removed by the **Flush** subcommand). All of the entries in the table, whether they are permanent or temporary, survive across switch reboots. Therefore, you must use the **Delete** subcommand when you want to remove permanent entries from the table.

5. The following prompt displays:

Use trailer encapsulation on this host (y/n)? (n) :

Enter **y** if you want to use trailer encapsulation on this host. Enter **n** if you do not want to use trailer encapsulation on this host.

6. The system will then confirm the addition to the table (an example is shown below).

ARP table entry for host 198.206.184.35 successfully added

7. The **xlat** submenu will then be redisplayed:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Deleting Entries from the ARP Table

The **Delete** subcommand is used to delete a “permanent” IP address from the ARP Table. Follow the steps below to delete an address from the ARP Table.

1. Enter **delete**.

The following prompt displays:

Host name or IP addr to delete:

Enter the host name or address that you wish to delete.

2. The system will then confirm the deletion from the table (an example is shown below).

ARP table entry for host 198.206.184.35 successfully deleted

3. The **xlat** submenu will then be redisplayed:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Flushing Temporary Entries from the ARP Table

The **Flush** subcommand is used to delete “temporary” IP addresses from the ARP Table. Follow the steps below to flush all temporary addresses from the ARP Table.

1. Enter **flush**.

The following prompt displays:

Flushing all non-permanent ARP table entries...done

2. The **xlat** submenu will then be redisplayed:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Finding a Specific IP Address in the ARP Table

The **Macfind** subcommand is used to locate a specific IP address in the ARP Table *based on knowing its MAC address*. (The **Ipfind** subcommand, discussed next, is used to find a specific MAC address based on a known IP address).

Follow the steps below to display a specific IP address in the ARP Table.

1. Enter **macfind**.

The following prompt displays:

MAC address to find (format aa:bb:cc:dd:ee:ff):

2. Enter the known MAC address (for example, 00:05:02:c0:7f:11).

A prompt similar to the following displays which shows the IP address that is related to the MAC address you entered:

Corresponding IP address: 198.206.184.34

3. The **xlat** submenu will then be redisplayed:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Finding a Specific MAC Address in the ARP Table

The **Ipfind** subcommand is used to locate a specific MAC address in the ARP Table *based on knowing its IP address or host name*. (The **Macfind** subcommand, discussed above, is used to find a specific IP address based on a known MAC address).

Follow the steps below to display a specific MAC address in the ARP Table.

1. Enter **ipfind**.

The following prompt displays:

Hostname or IP address to find:

2. Enter the known IP address or host name (for example, 198.206.184.34).

A prompt similar to the following displays which shows the MAC address that is related to the IP address you entered:

Corresponding MAC address: 00:05:02:c0:7f:11

3. The **xlat** submenu will then be redisplayed:

Enter command (Add/Delete/Show/Flush/Macfind/Ipfind/Quit) (Show) :

Viewing IP Statistics and Errors

The **ips** command is used to monitor IP datagram traffic and errors. The **ips** command displays *cumulative* IP statistics and errors. The statistics show the cumulative totals since the last time the switch was powered on or since the last reset of the switch was executed.

To display information about IP statistics and errors, enter the following command:

ips

The screen display depends on the type of switch you are using. For the OmniSwitch, Pizza-Switch, and OmniStack, a screen similar to the following displays:

IP Statistics and Errors

Default Time to Live	32
Reassembly Timeout (seconds)	1
Total Datagrams Recvd/Forwarded	77972 / 58177
HRE Datagrams Forwarded	0
PDU's Requested for Transmit	4294931545
PDU's Needing Reassembly	0
PDU's Successfully Reassembled	0
PDU's Needing Fragmentation	0
Fragments created	0
IP Errors (Discards due to the following problems)	
Header errors	0
Address errors	45994
Unknown/Unsupported Protocol	0
Local discards inbound/outbound	0 / 0
Unknown Route	45994
Reassembly Failures	0
Fragmentation Failures	0

The Omni Switch/Router (OmniS/R) includes fastpath code that enhances the speed of IP routing. The **fpx.img** file must be loaded on the OmniS/R for this feature to be enabled. Fast-path statistics are included on the IP Statistics and Errors screen. The screen display is similar to the following:

IP Statistics and Errors

Default Time to Live	32
Reassembly Timeout (seconds)	1
Total Datagrams Recvd/Forwarded	513342 / 513283
Fastpath Datagrams Received	513281
Fastpath Datagrams Forwarded	513280
Fastpath Inbound Discards	1
Fastpath Utilization	100%
PDU's Requested for Transmit	4294931545
PDU's Needing Reassembly	0
PDU's Successfully Reassembled	0
PDU's Needing Fragmentation	0
Fragments created	0
IP Errors (Discards due to the following problems)	
Header errors	0
Address errors	45994
Unknown/Unsupported Protocol	0
Local discards inbound/outbound	0 / 0
Unknown Route	45994
Reassembly Failures	0
Fragmentation Failures	0

The fields on this screen have the following meanings:

Default Time to Live

The default time, in seconds, assigned to each outgoing IP datagram before it is discarded as expired.

Reassembly Timeout (seconds)

The time, in seconds, to wait for all fragments to arrive before discarding datagrams.

Total Datagrams Recvd/Forwarded

The total number of input IP datagrams received, including those received in error.

HRE Datagrams Forwarded

The total number of IP datagrams forwarded by the HRE (Hardware Routing Engine).

Fastpath Datagrams Received

(Displays for Omni S/R only.) The number of IP datagrams received by the fastpath code.

Fastpath Datagrams Forwarded

(Displays for Omni S/R only.) The number of IP datagrams forwarded to their destination without using the MPM.

Fastpath Inbound Discards

(Displays for Omni S/R only.) The number of bad packets received and discarded. Typically this value should be zero.

Fastpath Utilization

(Displays for Omni S/R only.) The percentage of total datagrams received that are forwarded by the fastpath code.

PDU's Requested for Transmit

The total number of IP datagrams which transmit local IP user-protocols (including ICMP) supplied to IP in requests for transmission, not including forwarded datagrams.

PDU's Needing Reassembly

The number of IP datagram fragments that needed to be reassembled by this switch.

PDU's Successfully Reassembled

The number of IP datagrams successfully reassembled by this switch.

PDU's Needing Fragmentation

The number of IP datagrams requiring fragmentation by this switch.

Fragments created

The number of IP datagram fragments that have been generated as a result of fragmentation by this switch.

Header errors

The number of input IP datagrams discarded due to errors in their IP header, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discarded in processing their IP options, etc.

Address errors

The number of input IP datagrams discarded because the IP header destination field contained an invalid address.

Unknown/Unsupported Protocol

The number of local addresses, unsupported protocols, datagrams received successfully but discarded because of an unknown or unsupported protocol.

Local discards inbound/outbound

The number of packets discarded, both inbound and outbound, though they had no errors to prevent their being transmitted (lack of buffer space, etc.).

Unknown Route

The number of packets received and discarded by IP because IP was unable to route them.

Reassembly Failures

The number of failures detected by the IP reassembly algorithm for all reasons (timed out, error, etc.) This value is not necessarily a count of the discarded fragments.

Fragmentation Failures

The number of IP datagrams discarded because they needed to be fragmented but could not be. This situation could happen when a large packet has the "Don't Fragment" flag set.

Viewing the IP Forwarding Table

The **ipr** command is used to display the IP Forwarding Table. The entries in the table show the routes entered by a routing protocol, if the switch is running any of the supported protocols, and the static routes that you may have entered manually. You can also add to, or remove static routes from, the IP Forwarding Table (see *Adding an IP Static Route* on page 31-16 and *Removing an IP Static Route* on page 31-17).

To display the IP Forwarding Table, enter the following command:

```
ipr
```

A screen similar to the following displays:

10 routes in forwarding table

IP FORWARDING TABLE

Network	Mask	Gateway	Metric	Group VLAN Id:	Protocol
0.0.0.0	255.0.0.0	198.206.184.254	1	1:1	STATIC
10.0.0.0	255.0.0.0	10.0.0.1	1	6:1	STATIC
11.0.0.0	255.0.0.0	11.0.0.1	1	5:1	STATIC
90.0.0.0	255.0.0.0	90.0.0.3	1	4:1	DIRECT
127.0.0.0	255.0.0.0	127.0.0.1	0	1:2	LOOPBACK
127.0.0.1	255.255.255.255	127.0.0.1	0	2:3	LOOPBACK
196.196.7.0	255.255.255.0	196.196.7.42	1	3:1	RIP
198.206.187.0	255.255.255.0	198.206.183.0	1	1:4	STATIC
198.206.184.0	255.255.255.0	198.206.184.42	1	1:1	RIP
203.229.229.0	255.255.255.0	203.229.229.250	1	2:1	RIP

The fields on this screen have the following meanings:

Network

The destination network IP address.

Mask

The IP subnet mask.

Gateway

The network address of the gateway (the router from which this address was learned).

Metric

The metric associated with this network. Generally, this is a RIP “hop” count, or the number of hops the network is away from this router.

Group VLAN Id

The group and VLAN number from which this IP address was learned.

Protocol

The way in which this route was learned, for example, through RIP.

Adding an IP Static Route

The **aisr** command is used to add IP static routes to the switch's IP Forwarding Table. You might want to add a static route to send traffic to a router other than the one determined by the routing protocols.

In order to add a static route, you will need to know the host/net IP address and the gateway IP address which will be used to route traffic to the external IP address. The name "default" is used to specify the default route (the route used when no other routes are found).

Follow the steps below to add an IP static route.

1. Enter **aisr**.

A screen similar to the following displays:

Do you want to see the current route table? (y or n) (y) :

2. Enter a **y** at this prompt (or press **<Enter>**) to display the current routing table.

A screen similar to the following displays:

IP FORWARDING TABLE

Network	Mask	Gateway	Metric	Group VLAN Id:	Protocol
0.0.0.0	255.0.0.0	198.206.184.254	1	1:1	STATIC
10.0.0.0	255.0.0.0	10.0.0.1	1	6:1	STATIC
11.0.0.0	255.0.0.0	11.0.0.1	1	5:1	STATIC
90.0.0.0	255.0.0.0	90.0.0.3	1	4:1	DIRECT
127.0.0.0	255.0.0.0	127.0.0.1	0	1:2	LOOPBACK
127.0.0.1	255.255.255.255	127.0.0.1	0	2:3	LOOPBACK
196.196.7.0	255.255.255.0	196.196.7.42	1	3:1	RIP
198.206.187.0	255.255.255.0	198.206.183.0	1	1:4	STATIC
198.206.184.0	255.255.255.0	198.206.184.42	1	1:1	RIP
203.229.229.0	255.255.255.0	203.229.229.250	1	2:1	RIP

Destination IP address of host or network :

Enter the IP address of the host or network to which you are setting up a route. For a "default" route, use an entry of 0.0.0.0 as the IP address (or just enter the word **default**).

3. If you entered an IP address, a prompt similar to this one will be displayed:

Host or network mask (255.255.255.000) :

Enter the mask (or just press **<Enter>** to accept the default mask).

4. The following prompt displays:

IP address of next hop :

Enter the IP address of the next hop (the gateway) router to the destination IP address. The gateway address must be on the same network as one of the VLANs (that is, it must be a directly connected network).

A message will confirm the creation of the static route:

Route successfully added

Removing an IP Static Route

The **risr** command is used to remove IP static routes from the switch's IP Forwarding Table. Follow the steps below to remove an IP static route.

1. Enter **risr**.

A screen similar to the following displays:

Do you want to see the current route table? (y or n) (y) :

2. Enter a **y** at this prompt (or just press **<Enter>**) to display the current forwarding table.

A screen similar to the following displays:

IP FORWARDING TABLE

Network	Mask	Gateway	Metric	Group VLAN Id:	Protocol
0.0.0.0	255.0.0.0	198.206.184.254	1	1:1	STATIC
10.0.0.0	255.0.0.0	10.0.0.1	1	6:1	STATIC
11.0.0.0	255.0.0.0	11.0.0.1	1	5:1	STATIC
90.0.0.0	255.0.0.0	90.0.0.3	1	4:1	DIRECT
127.0.0.0	255.0.0.0	127.0.0.1	0	1:2	LOOPBACK
127.0.0.1	255.255.255.255	127.0.0.1	0	2:3	LOOPBACK
196.196.7.0	255.255.255.0	196.196.7.42	1	3:1	RIP
198.206.187.0	255.255.255.0	198.206.183.0	1	1:4	STATIC
198.206.184.0	255.255.255.0	198.206.184.42	1	1:1	RIP
203.229.229.0	255.255.255.0	203.229.229.250	1	2:1	RIP

Destination IP address of host or network :

Enter the IP address of the host or network that you want to remove.

3. A prompt similar to the following one displays:

Host or network mask (255.255.255.000) :

Enter the mask (or just press **<Enter>** to accept the default mask).

4. The following prompt displays:

IP address of next hop :

Enter the IP address of the next hop (the gateway) router to the destination IP address.

A message will confirm the deletion of the static route:

Route successfully deleted

Viewing ICMP Statistics and Errors

The **icmps** command is used to monitor ICMP activity.

To display information about ICMP statistics and errors, enter the following command:

icmps

A screen similar to the following displays:

ICMP Statistics		
	In	Out
Total ICMP Messages	1	1
Redirect Messages	0	0
Echo Messages	1	0
Echo Reply Messages	0	1
Time Stamp Messages	0	0
Time Stamp Reply Messages	0	0
Address Mask Messages	0	0
Address Mask Reply Messages	0	0
ICMP Errors		
	In	Out
Errors	0	0
Destination Unreachable Msgs	0	0
Time Exceeded Msgs	0	0
Parameter Problems	0	0
Source Quenches	0	0

The following field descriptions pertain to both the “in” and “out” statistics:

Total ICMP Messages

The total number of ICMP messages which this switch received or attempted to send out.

Redirect Messages

The number of ICMP Redirect messages sent/received by this switch.

Echo Messages

The number of ICMP Echo messages sent/received by this switch to see if a destination is active and reachable.

Echo Reply Messages

The number of ICMP Echo Reply messages received by this switch.

Time Stamp Messages

The number of Time Stamp Request messages sent/received by this switch requesting/receiving a reply with timestamp.

Time Stamp Reply Messages

The number of Time Stamp Reply messages sent/received by this switch.

Address Mask Messages

The number of Address Mask Reply messages that were sent/received by this switch in an attempt to determine the subnet mask for a network.

Address Mask Reply Messages

The number of Address Mask Reply messages that were sent/received by this switch.

Errors

The number of ICMP messages this switch sent/received but was unable to process because something was wrong (for example, a checksum failure).

Destination Unreachable Msgs

The number of ICMP “destination unreachable” messages that were sent/received. These occur when the gateway is unable to route a datagram to its destination.

Time Exceeded Msgs

The number of “time exceeded” messages that were sent/received. These occur when a packet is dropped because the Time-to-Live counter reaches zero. When a large number of these messages are encountered this is a symptom that packets are looping, that congestion is severe, or that the Time-to-Live counter is set too low. These messages also occur when all the fragments trying to be reassembled don't arrive before the reassembly timer expires.

Parameter Problems

The number of messages sent/received which indicate that an illegal value has been detected in a header field. These messages can indicate a problem in the sending host's IP software or possibly in the gateway's software.

Source Quenches

The number of messages sent/received which tell a host that is sending too many packets. A host should attempt to reduce its transmissions upon receiving these messages.

Using the PING Command

The **ping** command is used to test the reachability of IP network destinations. A fast ping command (**fping**) is also available for repeating the last ping request sent from the switch. The commands send an ICMP echo request to a destination and then wait for a reply.

Follow the steps below to issue an IP ping request.

1. Enter **ping**.

A screen similar to the following displays:

Host () :

Enter the IP address of the host that you wish to “ping.”

2. The following prompt displays:

Count (0 for infinite) (1) :

Enter the number of frames to be transmitted (0 equals “infinite”). To abort an “infinite” transmission once it is in progress, just press **Enter** again.

3. The following prompt displays:

Size (64) :

Enter the desired size of the data portion of the packet. You can specify a packet size or a range of packet sizes up to 8148. If you give a range, the switch will increment the packet size by 1 each time up to the top of the range. It will then wrap and continue from the bottom size of the range again until the total number of frames specified in the count has been sent. You can also set the increment by which the packet size is increased each time by entering a comma and an increment number after the size. For example, an entry of

1-100,5

will send out the number of frames specified in the “Count” prompt, starting with a frame size of 1 and incrementing up to a frame size of 100 in steps of 5. Be aware that if the “Count” is too small, the 100-byte frame size may never be reached. And, if the count is large enough the packet size will wrap and go back to 1.

4. The following prompt displays:

Timeout (1) :

Enter the number of seconds the program is to wait for a response before timing out.

5. After answering the previous prompt, a screen similar to the following displays:

**Ping starting, hit <RETURN> to stop
PING 198.206.184.18: 64 data bytes**

[0] .

**----198.206.184.18 PING Statistics----
1 packets transmitted, 1 packets received, 0% packet loss**

To repeat the last ping request, at the system prompt, enter the following command:

fping

The last ping issued on the switch is immediately sent again. If no ping was previously issued, a prompt for the host address displays and defaults are used for Count, Size, and Timeout.

Viewing UDP Statistics and Errors

The **udps** command is used to display a listing of UDP statistics and errors. The **udps** command displays cumulative statistics since the last time the switch was powered on or since the last reset of the switch was executed.

To display information about UDP statistics and errors, enter the following command:

udps

A screen similar to the following displays:

Total UDP datagrams received	:	831
Total UDP datagrams transmitted	:	22
Total Datagrams received w/unknown applications	:	0
Total UDP datagrams w/other Errors	:	0

The fields on this screen have the following meanings:

Total UDP datagrams received

The total number of UDP datagrams delivered to UDP applications.

Total UDP datagrams transmitted

The total number of UDP datagrams sent from this switch.

Total UDP datagrams received w/unknown applications

The total number of datagrams for which there was no application at the destination.

Total UDP datagrams w/other Errors

The total number of UDP datagrams that could not be delivered for reasons other than lack of application at the destination.

Viewing the UDP Listener Table

The **udpl** command is used to display the UDP Listener Table. This table contains information about the switch's UDP end-points on which a local application is currently accepting datagrams. The UDP Listener Table shows the local IP addresses for each UDP listener and the local port number for this listener. An IP address of zero (0.0.0.0) indicates that it is listening on all interfaces.

To view the UDP Listener Table, enter the following command:

udpl

A screen similar to the following appears:

UDP Listener Table				
Local Address/Port			Recv-Q	Send-Q
0.0.0.0	/	162	0	0
0.0.0.0	/	161	0	0
0.0.0.0	/	520	0	0
0.0.0.0	/	1024	0	0

Local Address/Port

The local IP address, and the local port number, for this UDP connection. In the case of a connection in the listen state, which is willing to accept connections for any IP interface associated with the node, the value 0.0.0.0 is used.

Recv-Q and Send-Q

For the SNMP Traps (port 162) this is the number transmitted (there is no receive).

For the SNMP Requests (port 161) this is the number of Request PDUs sent and the number of Response PDUs received.

For RIP (port 520) this is the number of packets received and transmitted.

Viewing RIP Statistics and Errors

The **rips** command is used to display RIP statistics and errors. This command displays cumulative statistics since the last time the switch was powered on, or since the last reset of the switch was executed.

To display information about RIP statistics and errors, enter the following command:

```
rips
```

A screen similar to the following displays:

```

                RIP Statistics
Rtr (Group ID:VLAN ID 1:1) IP Address 198.206.182.115 RIP Mode silent
In          4769          Out          0
Transmit Error    0      Non-zero field    0
Bad Version      0      Bad Metric      0
Bad Family       0      Bad Size       0
Bad Address      0      Bad Command    0

```

The fields on this screen have the following meanings:

In/Out

The total number of RIP packets received and transmitted on a per-virtual-LAN basis.

Transmit Error

The total number of RIP packets that were unable to be sent.

Bad Version

The total number of RIP messages delivered to the switch that were not version 1.

Bad Family

The number of packets received on this VLAN whose family ID was not of the Internet family.

Bad Address

The number of received packets whose IP address was not a Class A, B, or C.

Non-zero Field

The number of received packets whose mandated “must-be-zero” fields were not zero.

Bad Metric

The number of received packets with a routing entry’s metric that was out of range.

Bad Size

The number of received packets that were not compatible with the expected size.

Bad Command

The number of received packets whose command field was not a “request” or “response.”

Viewing TCP Statistics

The **tcps** command is used to monitor TCP traffic activity and check TCP configuration parameters. To reconfigure TCP parameters, see *Viewing the TCP Connection Table* on page 31-26.

To display information about TCP activity, enter the following command:

tcps

A screen similar to the following displays:

TCP Statistics

Round Trip Algorithm Used	:	RSRE (MIL-STD-1778)
Retransmission Min/Max Timeout	:	300/3000
Max Connections Allowed	:	Unlimited
Active Opens	:	76
Passive Opens	:	43
Attempt Fails	:	0
Established Resets	:	5
Currently Established	:	3
Total Segments Received	:	1117
Total Segments Sent	:	832
Total Segments Retransmitted	:	0
Total Segments Received w/err	:	0
Total Segments Sent w/RST flag	:	0

The fields on this screen have the following meanings:

Round Trip Algorithm Used

The algorithm used to determine the Timeout value used for retransmitting unacknowledged octets. The value is: RSRE (MIL-STD-1778).

Retransmission Min/Max Timeout

The minimum/maximum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds.

Max Connections Allowed

The maximum number of connections allowed. Currently, the number is unlimited.

Active Opens

The number of times TCP connections have made a direct transition to the “synSent” state from the “closed” state (refer to RFC 973).

Passive Opens

The number of times TCP connections have made a direct transition to the “synReceived” state from the “listen” state (refer to RFC 973).

Attempt Fails

The number of times TCP connections have made a direct transition to the “closed” state from either the “synSent” state or the “synReceived” state, plus the number of times TCP connections have made a direct transition to the “listen” state from the “synReceived” state.

Established Resets

The number of times TCP connections have made a direct transition to the “closed” state from either the “established” state or the “closeWait” state.

Currently Established

The number of TCP connections for which the current state is either “established” or “closeWait”.

Total Segments Received

The total number of segments received, including those received in error. This count includes segments received on currently established connections.

Total Segments Sent

The total number of segments sent, including those on current connections but excluding those containing only retransmitted octets.

Total Segments Retransmitted

The number of TCP segments transmitted containing one or more previously transmitted octets.

Total Segments Received w/err

The total number of TCP segments that are in error; for example, bad TCP checksums.

Total Segments Sent w/RST flag

The number of TCP segments containing the RST flag.

Viewing the TCP Connection Table

The **tcpc** command is used to check the current TCP connections available in the TCP Connection Table.

To display the TCP Connection Table, enter the following command:

tcpc

A screen similar to the following displays:

TCP Connection/Listener Table

Local Address/Port	Remote Address/Port	Recv-Q	Send-Q	Conn State
127.0.0.1 / 1090	27.0.0.1 / 1091	0	0	ESTABLISHED
127.0.0.1 / 1091	127.0.0.1 / 1090	0	322	ESTABLISHED
198.206.184.42 / 23	198.206.184.34 / 2057	0	0	ESTABLISHED
0.0.0.0 / 23	0.0.0.0 / 0	0	0	LISTEN
0.0.0.0 / 21	0.0.0.0 / 0	0	0	LISTEN

The fields on this screen have the following meanings:

Local Address/Port

The local IP address for this TCP connection/ the local port for this TCP connection. In the case of a connection in the listen state which is willing to accept connections for any IP interface associated with the node, the value 0.0.0.0 is used.

Remote Address/Port

The remote IP address/the remote port number for this TCP connection.

Recv-Q

The number of segments received on this port.

Send-Q

The number of segments sent on this port.

Conn State

Describes the state of the TCP connection, as defined in RFC 973. Possible values are: closed, listen, synSent, synReceived, established, finWait1, finWait2, closeWait, lastAck, closing, time-Wait, and deleteTCB.

Using the TELNET Command

The **telnet** command is used to connect to another system. All of the standard TELNET commands are supported by the software in the switch.

To initiate a TELNET session, enter the following command:

```
telnet
```

A screen similar to the following displays:

```
telnet>
```

To see a listing of the TELNET commands, enter the following command

```
?
```

A screen similar to the following displays:

Commands may be abbreviated. Commands are:

close	close current connection
display	display operating parameters
mode	try to enter line or character mode ('mode ?' for more)
open	connect to a site
quit	exit telnet
send	transmit special characters ('send ?' for more)
set	set operating parameters ('set ?' for more)
unset	unset operating parameters ('unset ?' for more)
status	print status information
toggle	toggle operating parameters ('toggle ?' for more)
environ	change environment variables ('environ ?' for more)
?	print help information

Enter the desired commands to establish and conduct your TELNET session.

Tracing an IP Route

The **tracroute** command is used to find the IP route from the local switch to a specified IP address destination. This command displays the individual hops to the destinations well as some timing information. When using the **tracroute** command, you must enter the name of the destination as part of the command line.

As an example, we might want to trace the route to “corporate.com.” To do so, we would enter this command:

```
tracroute corporate.com
```

A screen similar to the following displays:

```
tracroute to corporate.com (198.206.185.7),30 hops max,40 byte packets  
1 branch-wan-gw.CORPORATE.COM (198.206.181.252) 16 ms 0 ms 16 ms  
2 10.254.1.253 (10.254.1.253) 98 ms 81 ms 98 ms  
3 198.206.185.7 (198.206.185.7) 121 ms 81 ms 98 ms
```

Each number listed in this screen display corresponds to an individual hop. The time needed to reach that hop is shown (in milli-seconds) after the hop’s IP address. The time may be followed by one of the following codes:

- !** The TTL of the received ICMP message is less than or equal to 1.
- !H** The host was unreachable.
- !N** The network was unreachable.
- !P** The protocol was unreachable.

If the time is replaced by an asterisk (*), no response was received from the host during the default 3-second timeout period.

Flushing the RIP Routing Tables

The **ripflush** command is used to flush all entries in the RIP Routing Table. All existing routes, with the *exception* of static and direct routes, are removed from the table by entry of the **ripflush** command.

To flush the RIP Routing Table, enter the following command:

ripflush

No message is displayed; the system prompt simply reappears.

Configuring IP RIP Filters

The **ipfilter** command is used to add or delete an IP RIP Output or Input filter. The IP RIP Filtering feature gives you a means of controlling the operation of the IP RIP protocol. By using IP RIP filters, you can minimize the number of entries that are put into the IP Routing Table as well as improve overall network performance by eliminating unnecessary traffic.

Two types of IP RIP filters are available:

- **RIP Input** filters control which IP networks are allowed into the switch's IP Routing Table whenever IP RIP updates are received.
- **RIP Output** filters control the list of IP networks that are included in the RIP Updates sent out by the switch on any interface. Thus, RIP Output filters effectively control which networks the router advertises in the RIP updates it generates.

Here are some example uses of IP RIP filters:

- RIP Input and Output filters can be used to isolate entire network segments (and/or routers) in order to make the network "appear" differently to the network's various segments.
- RIP Input and Output filters can be used to reduce the overall amount of WAN traffic that is needed to advertise routes that should not be used by a particular network segment.

♦ Important Note ♦

The IP RIP Filtering feature works *only* with the switch's standard RIP routing protocol. If you elect to use Alcatel's Advanced Routing feature (GateD) to provide RIP routing functionality in your switch, you will not be able to activate IP RIP Filtering.

Adding a "Global" IP RIP Filter

Follow the steps below to add a "global" IP RIP Output or Input filter.

1. Enter **ipfilter**.

A screen similar to the following displays:

Selecting global IP filter:

Add or delete entry {add(a), delete(d)} (a) :

Enter **a** (or just press **Enter**) to select to add a filter.

2. The following prompt displays:

**Filter type{RIP Output(ro),
RIP Input(ri)} (ro) :**

Enter **ro** (or just press **Enter**) to add a RIP Output filter. Enter **ri** to add a RIP Input filter.

3. The following prompt displays:

Filter action {block(b), allow(a)} (a) :

Enter **a** (or just press **Enter**) to set the filter action to "allow."

Enter **b** to set the filter action to "block."

- The following prompt displays:

IP address (default: all networks) :

Enter the IP address of the network that is to be allowed or blocked by the filter (or just press **Enter** to use the default of all networks). If you choose the default you will *not* be prompted for the network mask (as is shown in the next step).

- The following prompt displays:

IP network mask (default: 255.255.255.0) :

Enter the IP network mask of the network that is to be allowed or blocked by the filter (or just press **Enter** to use the default mask of 255.255.255.0). Note that the default mask will vary depending on the class of the IP address you entered above.

- A message will appear indicating that the filter was successfully added:

ipfilter successfully added

Adding an IP RIP Filter For a Specific Group or VLAN

Follow the steps below to add an IP RIP Output or Input filter for a specific Group or VLAN.

- Enter the Group and VLAN numbers after the command like this: **ipfilter 1:1**.

A screen similar to the following displays:

Selecting IP filter for interface 1:1 :

Add or delete entry {add(a), delete(d)} (a) :

Filter action {block(b), allow(a)} (a) :

IP address (default: all networks) :

IP network mask (default: 255.255.255.0) :

ipfilter successfully added

Enter **a** (or just press **Enter**) to select to add a filter.

- The following prompt displays:

**Filter type{RIP Output(ro),
RIP Input(ri)} (ro) :**

Enter **ro** (or just press **Enter**) to add a RIP Output filter. Enter **ri** to add a RIP Input filter.

- The following prompt displays:

Filter action {block(b), allow(a)} (a) :

Enter **a** (or just press **Enter**) to set the filter action to “allow.” Enter **b** to set the filter action to “block.”

- The following prompt displays:

IP address (default: all networks) :

Enter the IP address of the network that is to be allowed or blocked by the filter (or just press **Enter** to use the default of all networks). If you choose the default you will *not* be prompted for the network mask (as is shown in the next step).

5. The following prompt displays:

IP network mask (default: 255.255.255.0) :

Enter the IP network mask of the network that is to be allowed or blocked by the filter (or just press **Enter** to use the default mask of 255.255.255.0). Note that the default mask will vary depending on the class of the IP address you entered above.

6. If the Group:VLAN is a WAN routing service, the following prompt displays:

Do you wish to apply this filter to a specific WAN endpoint? (n): y
Frame Relay VC or PPP Peer {vc(v), peer(p)} (v):

Enter **y** to apply this filter to a specific WAN endpoint.

7. The following prompt displays:

Frame Relay VC or PPP Peer {vc(v), peer(p)} (v):

Enter **v** (or just press **Enter**) to apply this filter to a Frame Relay VC.

Enter **p** if you want to apply this filter to a PPP Peer.

8. If you choose to apply the filter to a Frame Relay VC, this prompt will appear:

Slot/port:

Enter the slot and port numbers to which you want to apply this filter.

9. You will then be prompted for the virtual circuit (VC) to which to apply this filter:

VC:

Enter the VC to which you want to apply this filter.

10. If you choose to apply a filter to a PPP Peer, this one prompt will appear:

Peer ID:

Enter the Peer ID to which you want to apply this filter.

A message will appear indicating that the filter was successfully added.

IP RIP Filter Precedence

Whenever you use multiple “allow” filters you must first define a filter to block all RIPs. Then, all of the seceding “allow” filters of the same type must be *at least* as specific in all areas in order for the filters to work. Note that filtering precedence is related only to “allow” filters. Multiple “block” filters can be defined with varying specificity in each of the areas of the filter. The filtering done by the configurable parameters (Address/Mask) in the “allow” filter must be at least as specific as the filtering defined in the “block” filter.

Deleting IP RIP Filters

Follow the steps below to delete an existing IP RIP Output or Input filter.

1. Enter **ipfilter**.

A screen similar to the following displays:

Selecting global IP filter:

Add or delete entry {add(a), delete(d)} (a) :

Enter **d** to select to delete a filter.

2. A screen similar to the following displays:

Displaying all filters:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP OUT	99.0.0.0	255.0.0.0	A	global
2	RIP IN	99.0.0.0	255.0.0.0	B	2:1

Entry number to delete? (default: none) :

This screen contains a list of the existing IP RIP filters. The fields on this screen are described in the next section (see *Displaying IP RIP Filters* on page 31-34).

3. Enter the index number of the filter that you want to delete. If you decide at this point that you want to abort out of the deletion process, simply press **Enter** to accept the default of “none”.
4. A message will confirm the deletion of the filter:

ipfilter successfully deleted

Displaying IP RIP Filters

The **ipf** command is used to display a list of all existing IP RIP Output and Input filters. See *Configuring IP RIP Filters* on page 31-30 for complete information on creating these filters.

Displaying a List of All IP RIP Filters

To display the listing of all existing IP RIP filters, enter the following command:

```
ipf
```

A screen similar to the following displays:

Displaying all filters:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP OUT	99.0.0.0	255.0.0.0	A	global
2	RIP IN	99.0.0.0	255.0.0.0	B	2:1
3	RIP OUT	All Networks		B	5:1 (3/1/32)
4	RIP IN	All Networks		B	6:1 (P1)

This screen contains a list of the existing IP RIP filters. The fields on this screen have the following meanings:

#

Indicates the index number assigned to identify this filter.

Type

Indicates the type of filter, either RIP Input (**RIP IN**) or RIP Output (**RIP OUT**).

Network

Indicates the IP address that is to be filtered (entered in dotted-decimal format). An entry of “All Networks” means that all addresses are to be filtered.

Mask

The IP network mask of the network to be filtered (entered in dotted-decimal format). This field is blank if the network entered is “All Networks.”

Md

Indicates the filter’s mode of operation, either to “allow” traffic (**A**) or to “block” traffic (**B**).

GP:VL (s/p/vc) or (Peer ID)

The first number (**GP**) is the Group associated with this entry. The second number (**VL**) is the VLAN associated with this entry. When a filter applies to all interfaces, this field will say “global.” If an entry refers to a Frame Relay interface, column headings for slot, port, and virtual circuit (**s/p/vc**) may be displayed when the filter is applied to a particular virtual circuit rather than to the entire VLAN. If an entry refers to a PPP interface, the Peer ID (**Peer ID**) may be displayed when the filter is applied to a particular PPP Peer.

Displaying a List of “Global” IP RIP Filters

To display a listing of just the global IP RIP filters, enter the following command:

```
ipf global
```

A screen similar to the following displays:

Displaying global filters:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP OUT	99.99.99.99	255.0.0.0	A	global

Displaying a List of Specific IP RIP Filters

To display a listing of IP RIP filters for a specific interface, you can specify other parameters along with the **ipf** command. The format for the command in this case is:

```
ipf <type> <GP:VL>
```

The type is one of these codes:

ri for RIP INput

ro for RIP OUTput

For example, to display a list of the filters defined for Group 2, VLAN 1, you would enter:

```
ipf 2:1
```

A screen similar to the following would be displayed:

Displaying filters for interface 2:1:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP IN	99.0.0.0	255.0.0.0	B	2:1

As another example, to display a list of all global RIP Output filters, you would enter:

```
ipf ro global
```

A screen similar to the following would be displayed:

#	Type	Network	Mask	Md	GP:VL (s/p/vc) (Peer ID)
1	RIP OUT	99.0.0.0	255.0.0.0	A	global

Viewing the IP-to-MAC Address Table

The **ipmac** command is used to display the IP-to-Mac Address Association Table. This table contains a listing of IP addresses and their associated MAC (Media Access Control) addresses together with the slot/port from which the information was learned. The information in this table is learned from ARP (Address Resolution Protocol) messages received on “leaf” ports. A “leaf” port is one on which Spanning Tree has been disabled or on which no Spanning Tree BPDUs have yet been received.

The **ipmac** command can be very helpful in resolving certain problems. For example, in large networks where hosts are frequently moved around, users can experience connectivity problems. In this situation, the **ipmac** command can be used to help locate a particular IP workstation. Another use is to help resolve duplicate IP addresses on a network. The program checks all ARP messages, whether they are received on a “leaf” port or not, against those in its table to see if a duplicate IP address exists. If a duplicate is detected, an SNMP trap message is generated and the duplicate can easily be seen in the table produced by the **ipmac** command.

The **ipmac** command can be entered alone in which case it will display all entries currently in the table, or you may enter a specific IP address along with the command to show only the information related to that IP address. An optional parameter (-f) can be entered to flush the table. Each of these uses of the **ipmac** command is illustrated below.

Displaying All Entries in the IP-to-MAC Table

To display the list of all the entries in the IP-to-MAC table, enter the following command:

```
ipmac
```

A screen similar to the following displays:

IP to MAC ADDRESS ASSOCIATION TABLE

IP Address	MAC Address	Slot / Intf
192. 168. 10. 1	0020DA:6DE610	4 / 5
172. 16. 0. 5	0020DA:76D3D0	3 / 2
172. 16. 0. 7	00E029:00D41E	3 / 2
172. 16. 0. 41	0000C0:24FFEC	3 / 2
172. 16. 0. 47	00A0C9:0AA907	3 / 2
172. 16. 0. 28	0020DA:7AE9D3	3 / 2
172. 16. 0. 45	080020:8AE301	3 / 2
172. 16. 0. 60	0020DA:73C3A0	3 / 2
172. 16. 30. 00	0020AF:04BA57	3 / 2
172. 16. 41. 03	0000C0:AD8EE9	3 / 2
172. 16. 50. 12	080020:7B79E1	3 / 2
172. 16. 255. 254	0020DA:6F97E5	3 / 2
*****	0020DA:032273	5 / 1
192. 168. 10. 1	0020DA:7AEA60	3 / 2
198. 206. 182. 222	0020DA:7F48A0	3 / 2

The fields on this screen have the following meanings:

IP Address

The IP address learned from ARP messages received on “leaf” ports. A series of asterisks (*****) in this field indicates that the preceding entry is a duplicate to this entry. In the example screen shown above, the address 172.16.255.254 is assigned to two MAC addresses.

MAC Address

The MAC address corresponding to the listed IP address.

Slot/Intf

The slot number and interface number from which the IP and MAC addresses were learned.

Displaying Information for a Specific IP Address

To display the entry in the IP-to-MAC table for a specific IP address, enter the desired IP address after the command. For example to located the entry for IP address 192.168.10.1, you would enter the following command:

```
ipmac 192.168.10.1
```

A screen similar to the following displays:

IP to MAC ADDRESS ASSOCIATION TABLE

IP Address	MAC Address	Slot / Intf
192.168. 10. 1	0020DA:6DE610	4 / 5

Flushing Entries from the Table

To flush all the entries in the IP-to-MAC table, enter the following command:

```
ipmac -f
```

No message will be displayed; the system prompt will reappear when the flush is completed.

Path MTU Discovery

All Gigabit Ethernet modules and all Mammoth-based Ethernet modules on the OmniSwitch and Omni Switch/Router in Release 4.0 and later support path Maximum Transmission Unit (MTU) discovery. In path MTU discovery, the Ethernet frame (datagram) size is set to the largest size that does not require fragmentation anywhere along the path from a source host to its destination. This frame size, known as a Path MTU (PMTU), is thus equal to the minimum of the MTUs of each hop in the path.

◆ Note ◆

MTU discovery is *not* supported on token ring, FDDI, WAN, or non-Mammoth Ethernet modules. However, token ring and FDDI can be used as intermediate links (e.g., trunking or bridging) between remote switches.

Path MTU discovery is active all of the time and is part of the switch's operating system; you do not need configure it.

The source host initially assumes that the PMTU of a path is the MTU of the first hop. It sends all datagrams with the "Don't Fragment" (DF) bit set. If a switch/router along the path receives a datagram that is too large to forward without fragmentation, the following steps will be executed:

1. The switch/router that cannot forward these datagrams (i.e., the constricting hop) will discard them.
2. The constricting hop will send ICMP destination unreachable messages to the source host with a code that indicates fragmentation is needed and the "Don't Fragment" (DF) bit in the Internet Protocol (IP) header has been set. This message (known as a "Datagram Too Big" message) contains the PMTU of the constricting hop.
3. After receiving a "Datagram Too Big" message, the source host reduces the size of the MTU so it matches the PMTU of the constricting hop.
4. The MTU discovery process ends when datagrams can be sent without fragmentation. However, the source host will *not* reduce the size of a datagram below 68 octets.